



Breaking Down the Entry Barriers on Linux Kernel Networking Stack

Juhee Kang(claudiajkang@gmail.com)





Speaker



Juhee Kang

Open-Source Part Frontier Developer

Project: Linux Kernel Networking Stack, Kubernetes i18n

Interest: Cloud, Network, Backend, etc...

Special Lecture at Universities

- Open-Source Software Development Mechanism.



INDEX

1

**Linux Kernel
Networking Stack**

2

Kernel Ecosystem

3

**Step by Step
Kernel Contribution**



Let's think about it

Linux Kernel Contribution?




Linux Kernel Contribution?



How do I start?



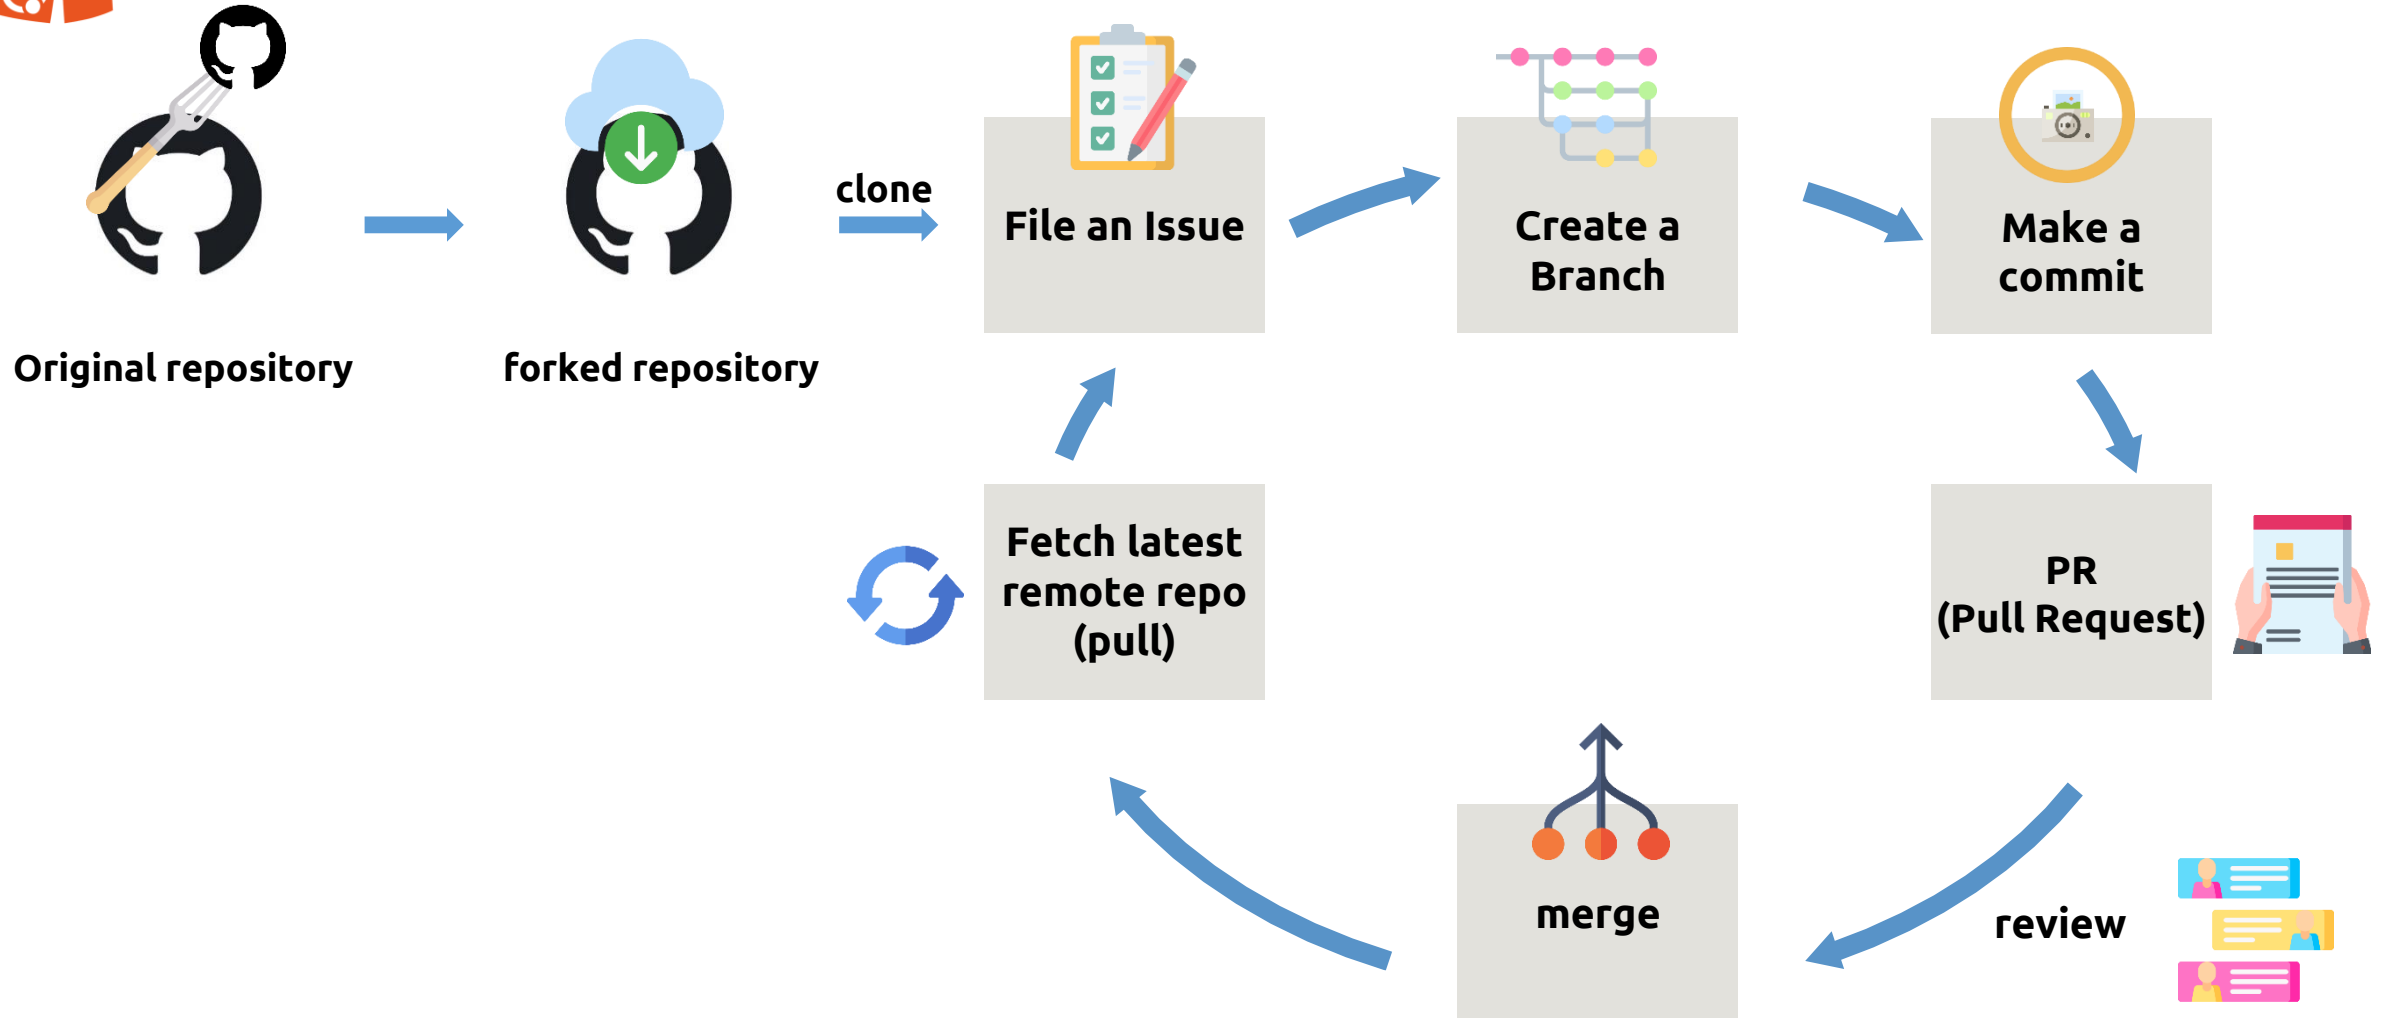
Contribution like other open-sources?



Maybe, contribute through GitHub?
(like any other open source)



Open-Source Contribution: managed with GitHub





Is Kernel Contribution done in the same way?



Is Kernel Contribution done in the same way?

torvalds / linux Public

Watch 8.1k

Fork 45.1k

Starred 140k

Code Pull requests 312 Actions Projects Security Insights

master 1 branch 763 tags

Go to file

Add file

Code



torvalds Merge tag 'for-6.1-rc1-tag' of git://git.kernel.org/pub/scm/linux/ker... aae703b 2 days ago 1,135,998 commits

Documentation	Merge tag 'random-6.1-rc1-for-linux' of git://git.kernel.org/pub/scm/...	4 days ago
LICENSES	LICENSES/LGPL-2.1: Add LGPL-2.1-or-later as valid identifiers	10 months ago
arch	Merge tag 'random-6.1-rc1-for-linux' of git://git.kernel.org/pub/scm/...	4 days ago
block	Merge tag 'random-6.1-rc1-for-linux' of git://git.kernel.org/pub/scm/...	4 days ago
certs	certs: make system keyring depend on built-in x509 parser	27 days ago
crypto	treewide: use get_random_bytes() when possible	9 days ago
drivers	Merge tag 'v6.1-p2' of git://git.kernel.org/pub/scm/linux/kernel/git/...	3 days ago
fs	Merge tag 'for-6.1-rc1-tag' of git://git.kernel.org/pub/scm/linux/ker...	2 days ago
include	Merge tag 'cgroup-for-6.1-rc1-fixes' of git://git.kernel.org/pub/scm/...	3 days ago
init	Merge tag 'mm-nonmm-stable-2022-10-11' of git://git.kernel.org/pub/s...	8 days ago
ipc	Merge tag 'mm-nonmm-stable-2022-10-11' of git://git.kernel.org/pub/s...	8 days ago

About

Linux kernel source tree

Readme

View license

140k stars

8.1k watching

45.1k forks

Releases

763 tags

Packages

No packages published

In torvalds/linux Kernel GitHub repository,



Is Kernel Contribution done in the same way?

torvalds / linux Public

Watch 8.1k

Fork 45.1k

Starred 140k

Code Pull requests 312 Actions Projects Security Insights

Filters is:pr is:open

Labels 0

Milestones 0

New pull request

312 Open ✓ 449 Closed

Author Label Projects Milestones Reviews Assignee Sort

Fix for packets being rejected in the ring buffer used by the xHCI co...

#805 opened on 2 Mar 2020 by TarAldarion

1

Update .clang-format

#803 opened on 29 Feb 2020 by farley-melo

1

Remove Intel unavailable hardware from SDIO IDs

#802 opened on 25 Feb 2020 by laserjobs

1

Improve readability of socket err

#801 opened on 18 Feb 2020 by sparshrestha

2

fix "unsigned" warn when running independently

#800 opened on 18 Feb 2020 by daTokenizer

1

Add support for USB modems using Conexant CX06836

#798 opened on 12 Feb 2020 by stevenyu1973

1

Pull Request exists just as other open-source.

#797 opened on 11 Feb 2020 by yunbotsai

1



Is Kernel Contribution done in the same way?

torvalds / linux Public

Watch 8.1k

Fork 45.1k

Starred 140k

Code Pull requests 312 Actions Projects Security Insights

Fix for packets being rejected in the ring buffer used by the xHCI co... #805

Code

Open TarAldarion wants to merge 1 commit into torvalds:master from TarAldarion:MRU_FIXES

Conversation 1 Commits 1 Checks 0 Files changed 1

+7 -0



TarAldarion commented on 2 Mar 2020

...ntroller. When a packet larger than MTU arrives in Linux from the modem, it is discarded with -EOVERFLOW error (Babble error). This is seen on USB3.0 and USB2.0 busses. This is essentially because the MRU (Max Receive Size) is not a separate entity to the MTU (Max Transmit Size) and the received packets can be larger than those transmitted. Following the babble error there were an endless supply of zero-length URBs which are rejected with -EPROTO (increasing the rx input error counter each time). This is only seen on USB3.0. These continue to come ad infinitum until the modem is shutdown. There appears to be a bug in the core USB handling code in Linux that doesn't deal well with network MTUs smaller than 1500 bytes. By default the dev->hard_mtu (the real MTU) is in lockstep with dev->rx_urb_size (essentially an MRU), and it's the latter that is causing trouble. This has nothing to do with the modems; the issue can be reproduced by getting a USB-Ethernet dongle, setting the MTU to 1430, and pinging with size greater than 1406.

Fix for packets being rejected in the ring buffer used by the xHCI co... 637f847

Reviewers

No reviews

Assignees

No one assigned

Labels

None yet

Projects

None yet



KernelPRBot commented

As KernelPRBot replied,

Hi @TarAldarion!



Is Kernel Contribution done in the same way?

Hi @TarAldarion!

Thanks for your contribution to the Linux kernel!

Linux kernel development happens on mailing lists, rather than on GitHub - this GitHub repository is a read-only mirror that isn't used for accepting contributions. So that your change can become part of Linux, please email it to us as a patch.

Sending patches isn't quite as simple as sending a pull request, but fortunately it is a well documented process.

Here's what to do:

- Format your contribution according to kernel requirements
- Decide who to send your contribution to
- Set up your system to send your contribution as an email
- Send your contribution and wait for feedback

How do I format my contribution?

The Linux kernel community is notoriously picky about how contributions are formatted and sent. Fortunately, they have documented their expectations.

Firstly, all contributions need to be formatted as patches. A patch is a plain text document showing the change you want to make to the code, and documenting why it is a good idea.

You can create patches with `git format-patch`.

Secondly, patches need 'commit messages', which is the human-friendly documentation explaining what the change is and why it's

Notifications

Customize

 [Subscribe](#)

You're not receiving notifications from this thread.

2 participants



**Linux Kernel development doesn't happen on GitHub.
Read-Only mirror purpose**



Is Kernel Contribution done in the same way?

Hi @TarAldarion!

Thanks for your contribution to the Linux kernel!

Linux kernel development happens on mailing lists, rather than on GitHub - this GitHub repository is a read-only mirror that isn't used for accepting contributions. So that your change can become part of Linux, please email it to us as a patch.

Sending patches isn't quite as simple as sending a pull request, but fortunately it is a well documented process.

Here's what to do:

- Format your contribution according to kernel requirements
- Decide who to send your contribution to
- Set up your system to send your contribution as an email
- Send your contribution and wait for feedback

How do I format my contribution?

The Linux kernel community is notoriously picky about how contributions are formatted and sent. Fortunately, they have documented their expectations.

Firstly, all contributions need to be formatted as patches. A patch is a plain text document showing the change you want to make to the code, and documenting why it is a good idea.

You can create patches with `git format-patch`.

Secondly, patches need 'commit messages', which is the human-friendly documentation explaining what the change is and why it's necessary.

need to comply with.

detailed guide for kernel contribution

Notifications

Customize

 **Subscribe**

You're not receiving notifications from this thread.

2 participants





Kernel Contribution?



**Instead of using GitHub,
Linux kernel uses its own git service.**

- Linux Kernel

<https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git>

The screenshot shows the web interface for the Linux kernel source tree repository. The browser address bar shows the URL `git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git`. The page title is "index : kernel/git/torvalds/linux.git" with the subtitle "Linux kernel source tree". Navigation tabs include "about", "summary", "refs", "log", "tree", "commit", "diff", and "stats".

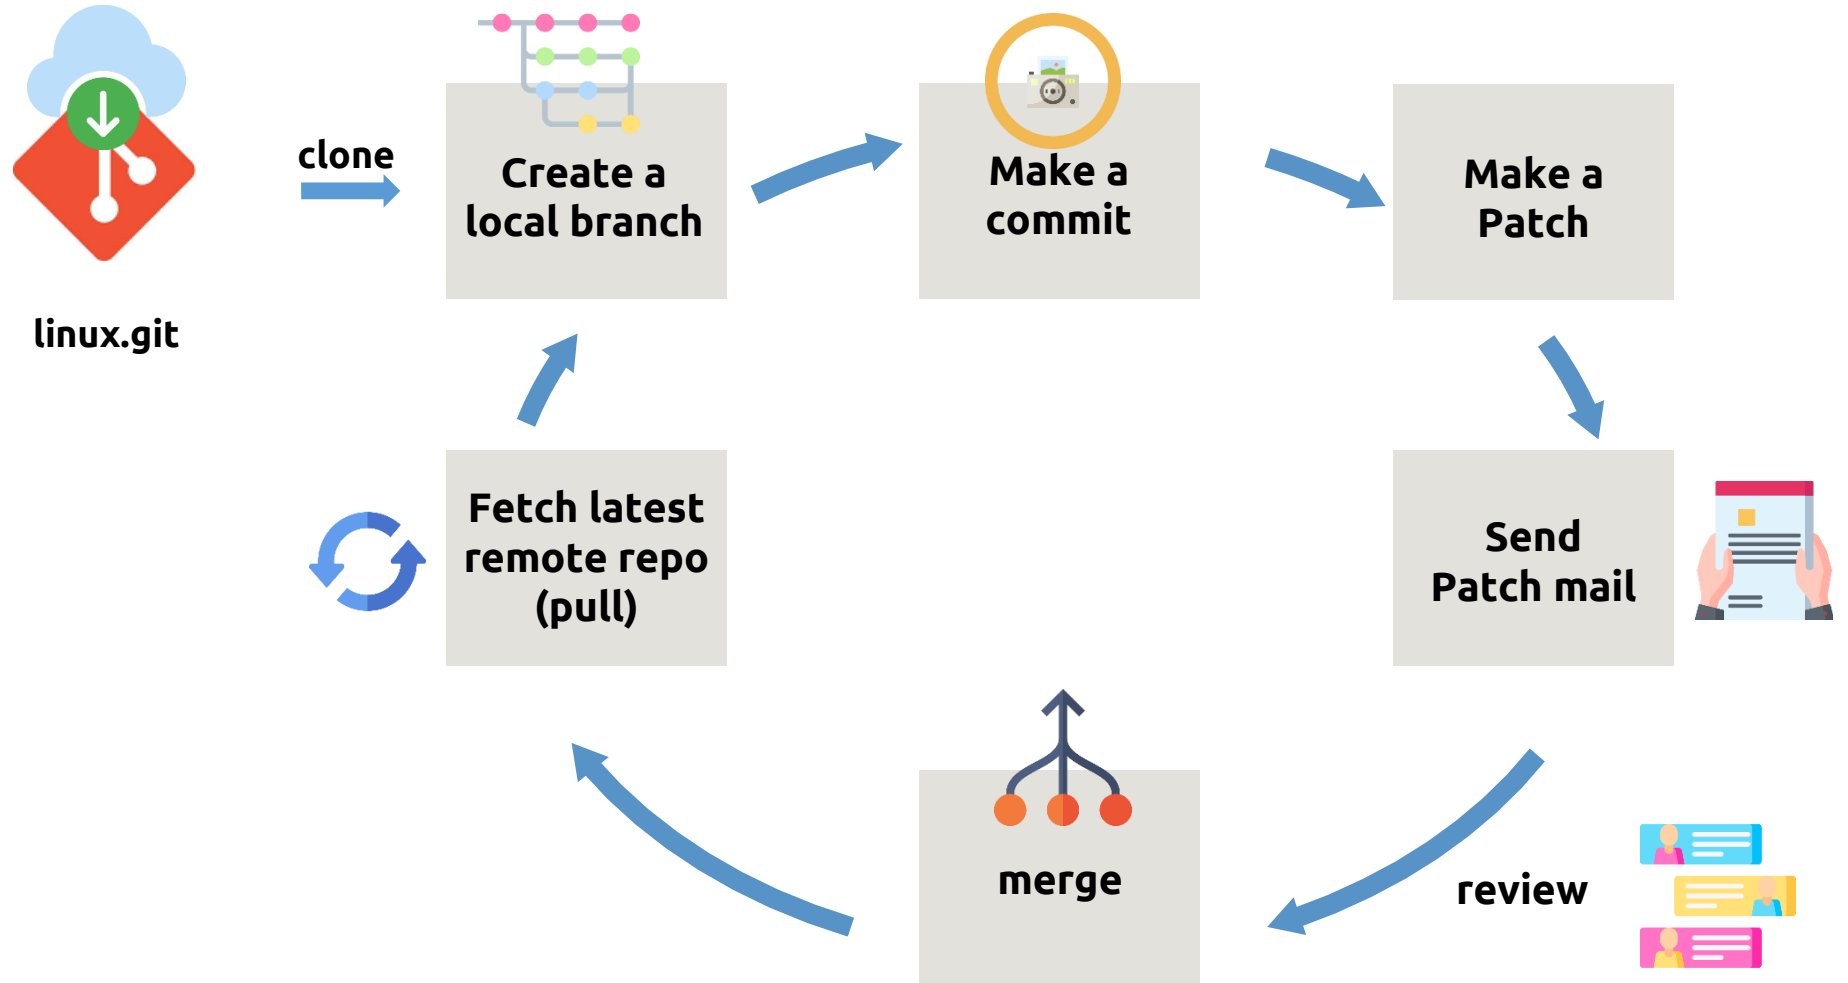
Branch	Commit message
master	Merge tag 'loongarch-fixes-5.19-4' of git://git.kernel.org/pub/scm/linux/kern...

Tag	Download
v5.19-rc5	linux-5.19-rc5.tar.gz
v5.19-rc4	linux-5.19-rc4.tar.gz
v5.19-rc3	linux-5.19-rc3.tar.gz
v5.19-rc2	linux-5.19-rc2.tar.gz
v5.19-rc1	linux-5.19-rc1.tar.gz
v5.18	linux-5.18.tar.gz
v5.18-rc7	linux-5.18-rc7.tar.gz
v5.18-rc6	linux-5.18-rc6.tar.gz
v5.18-rc5	linux-5.18-rc5.tar.gz
v5.18-rc4	linux-5.18-rc4.tar.gz
[...]	

Age	Commit message
8 hours	Merge tag 'loongarch-fixes-5.19-4' of git://git.kernel.org/pub/scm/linux/kern... HEAD master
9 hours	Merge tag 'net-5.19-rc6' of git://git.kernel.org/pub/scm/linux/kernel/git/net...
9 hours	Merge tag 'pinctrl-v5.19-2' of git://git.kernel.org/pub/scm/linux/kernel/git/...
9 hours	signal handling: don't use BUG_ON() for debugging
16 hours	LoongArch: Fix section mismatch warning
16 hours	LoongArch: Fix build errors for tinyconfig
16 hours	LoongArch: Remove obsolete mentions of vcsr
16 hours	LoongArch: Drop these obsolete selects in Kconfig
23 hours	Merge branch 'wireguard-patches-for-5-19-rc6'
23 hours	wireguard: Kconfig: select CRYPTO_CHACHA_S390
[...]	

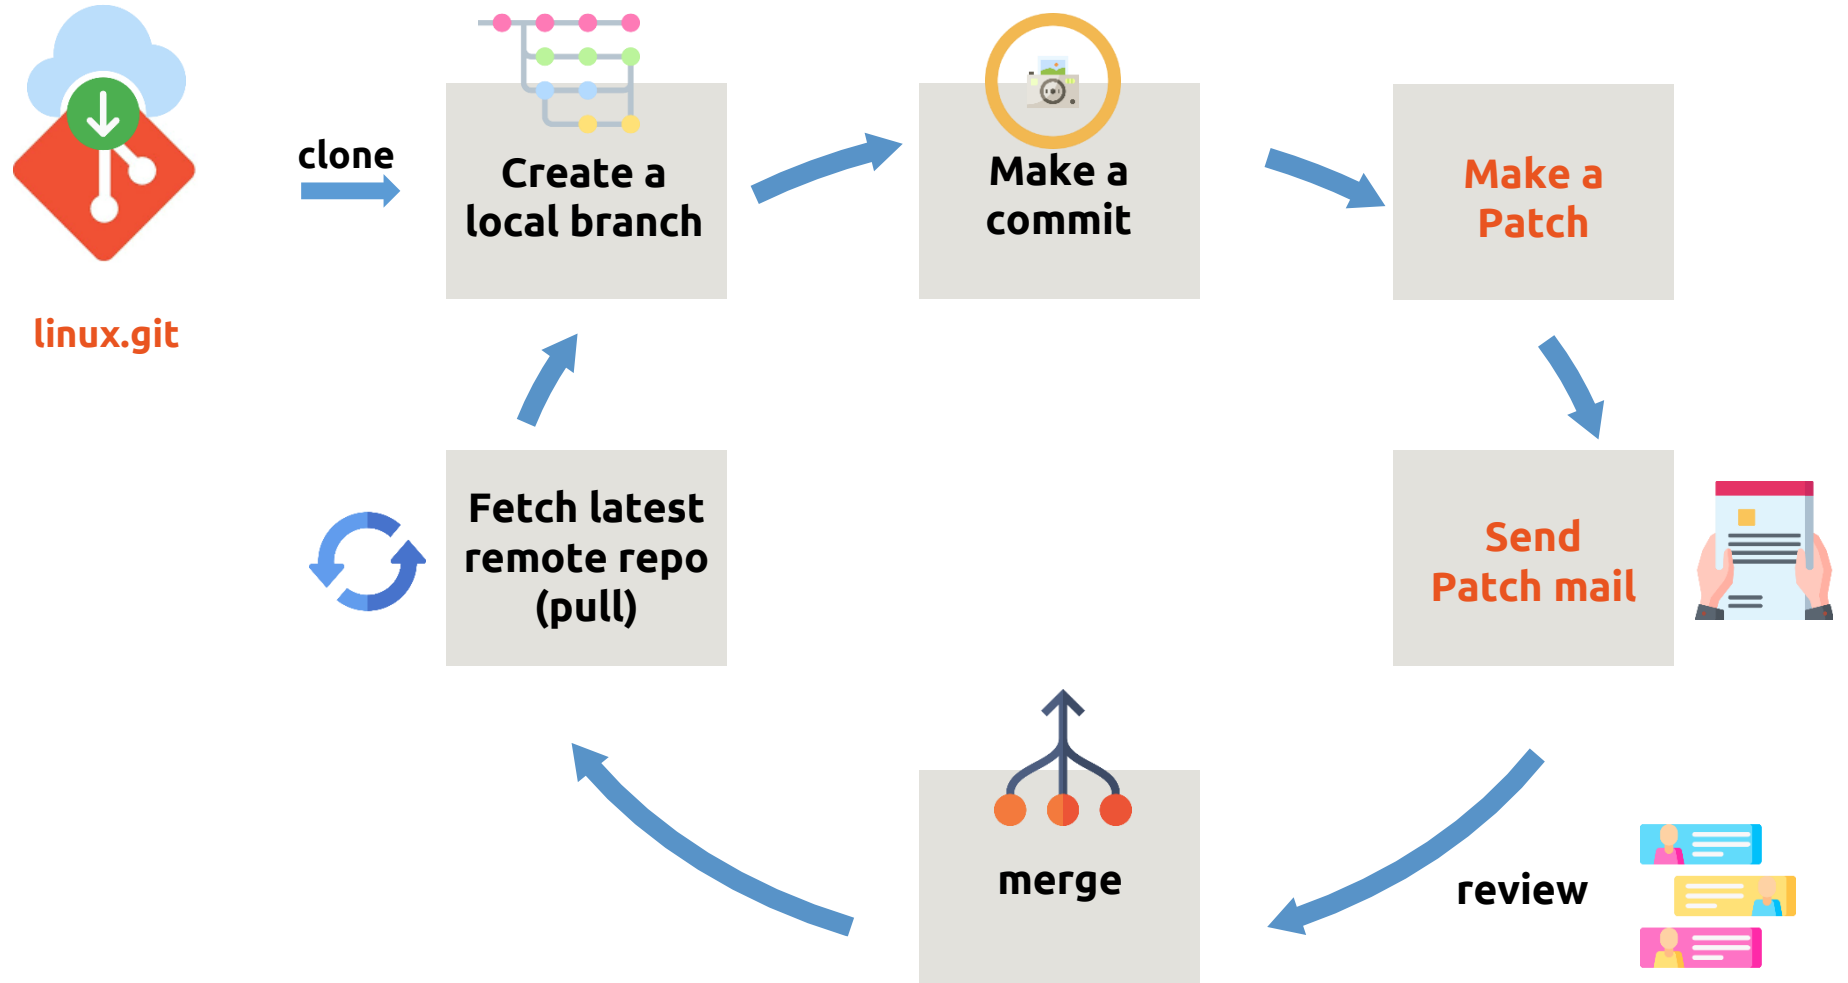


Kernel Contribution process?





Kernel Contribution process?






Kernel Contribution

Where to get source code?

Committing by email?



What is a patch?

How to configure build env?

Countless barriers before we even started...



Goal of this session

```
$ cd samples/bpf
$ make -j
$ ./tracex7
sh: 1: Syntax error: word unexpected (expecting ")")
```

In this scenario, suppose we fix a bug.



Goal of this session

Step through the Linux kernel contribution to break down contribution barriers.




1

Linux Kernel Networking Stack



Linux Kernel Networking Stack


Ethernet Frame






Linux Kernel Networking Stack

Ethernet
Frame




Network
Hardware







Linux Kernel Networking Stack

Ethernet
Frame




Network
Hardware







Linux Kernel Networking Stack

Ethernet
Frame



Network
Hardware

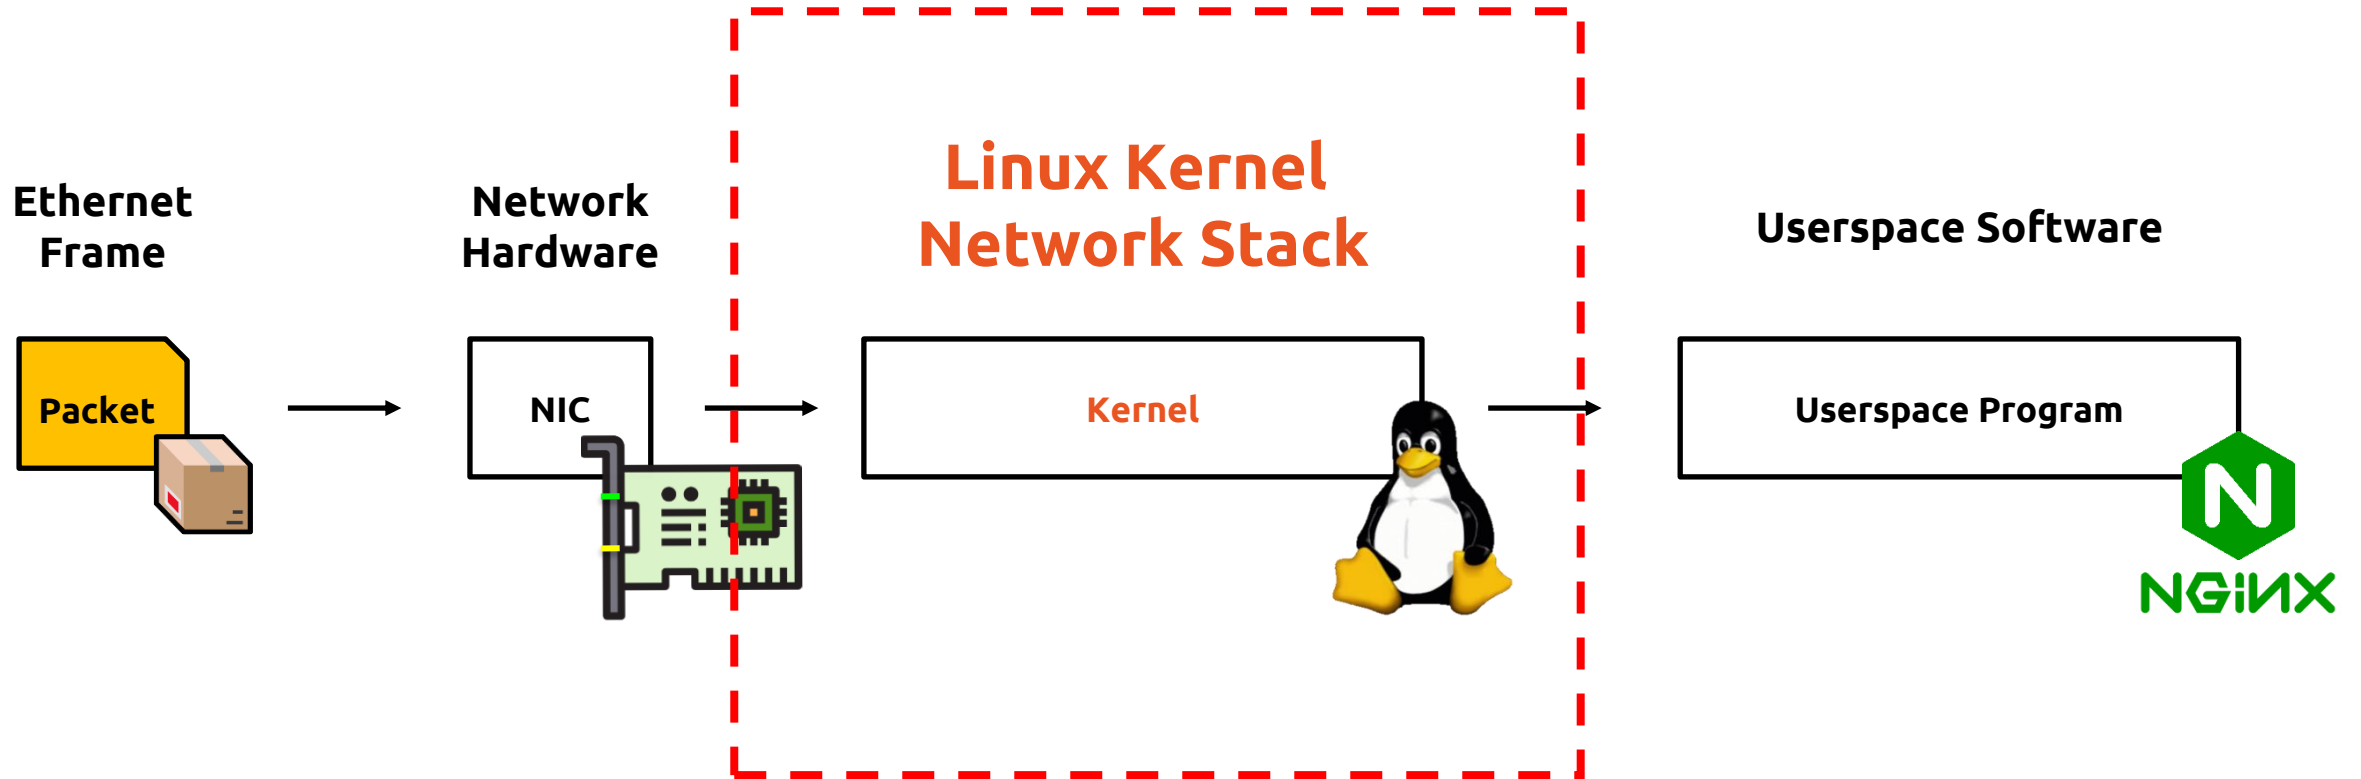


Userspace Software



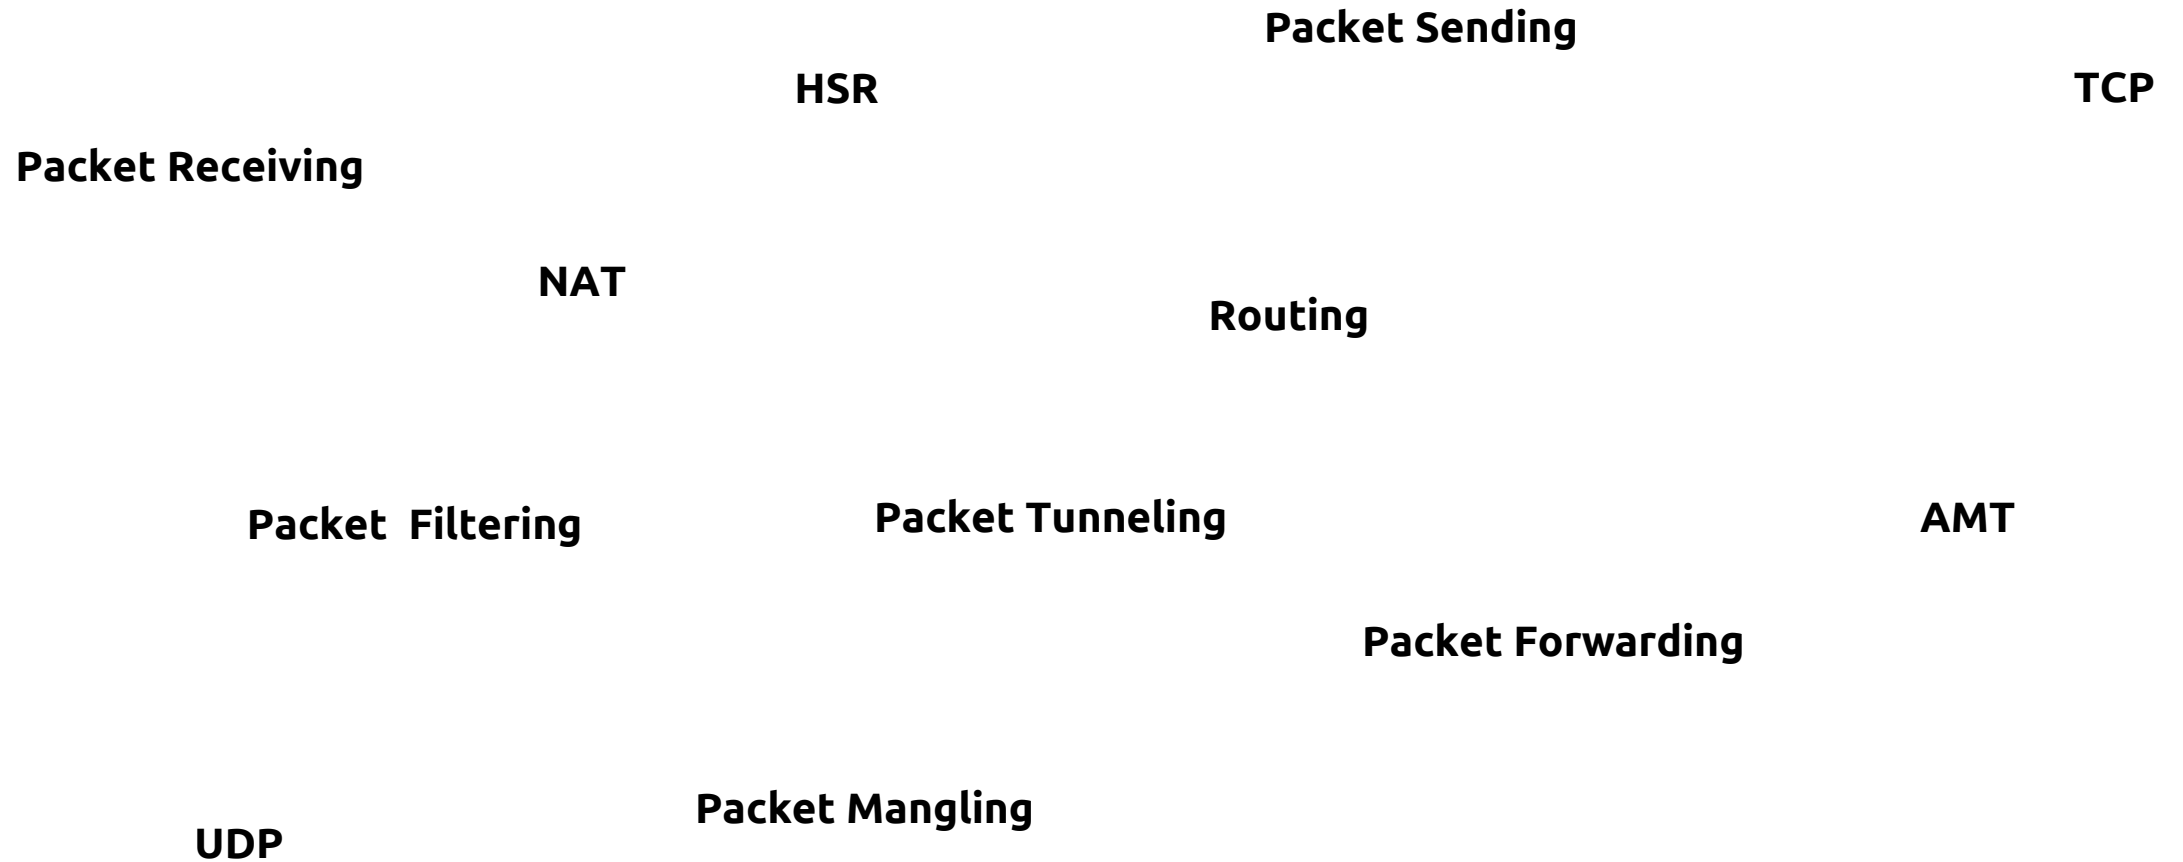


Linux Kernel Networking Stack



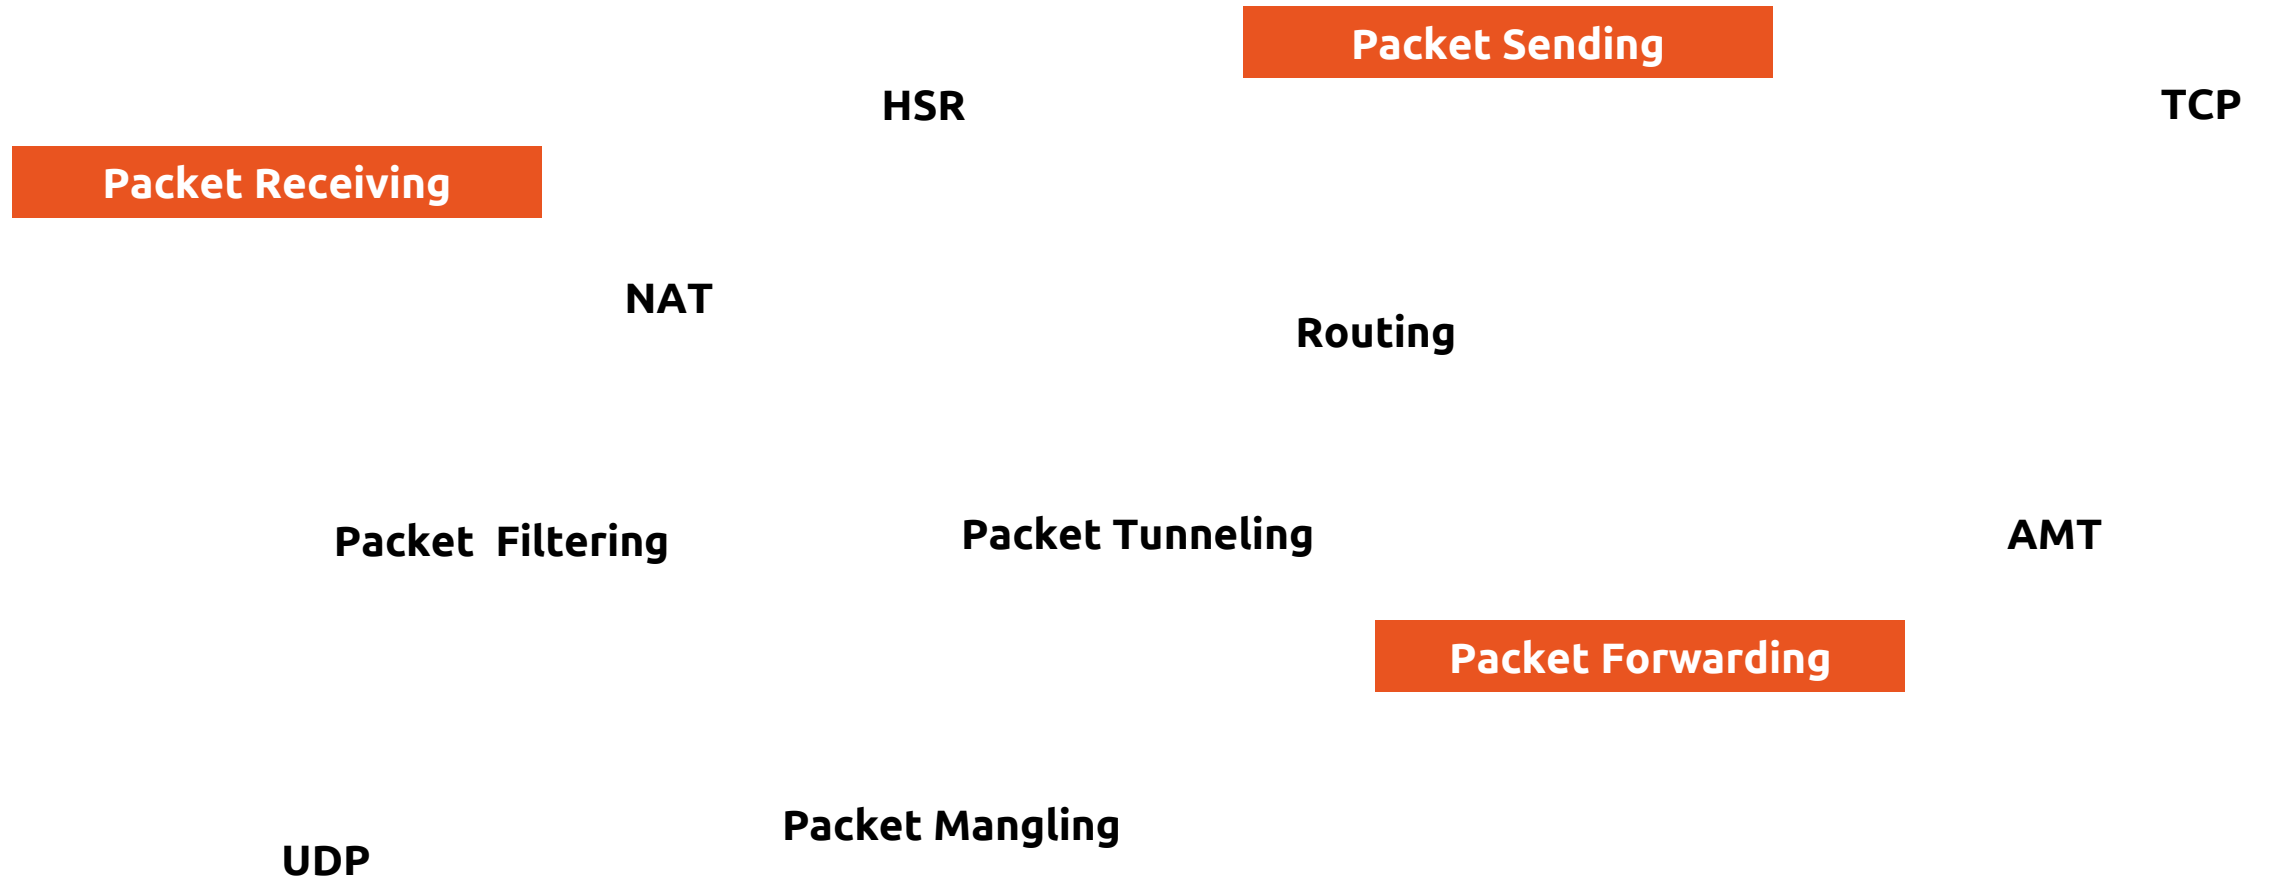


Linux Kernel Networking Stack





Basic Packet Path



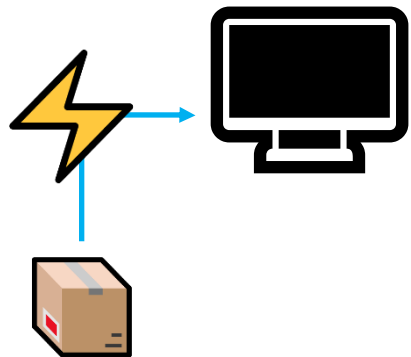


Basic Packet Rx/Tx Path




Packet Path


① Receiving Path



② Forwarding Path




③ Sending Path





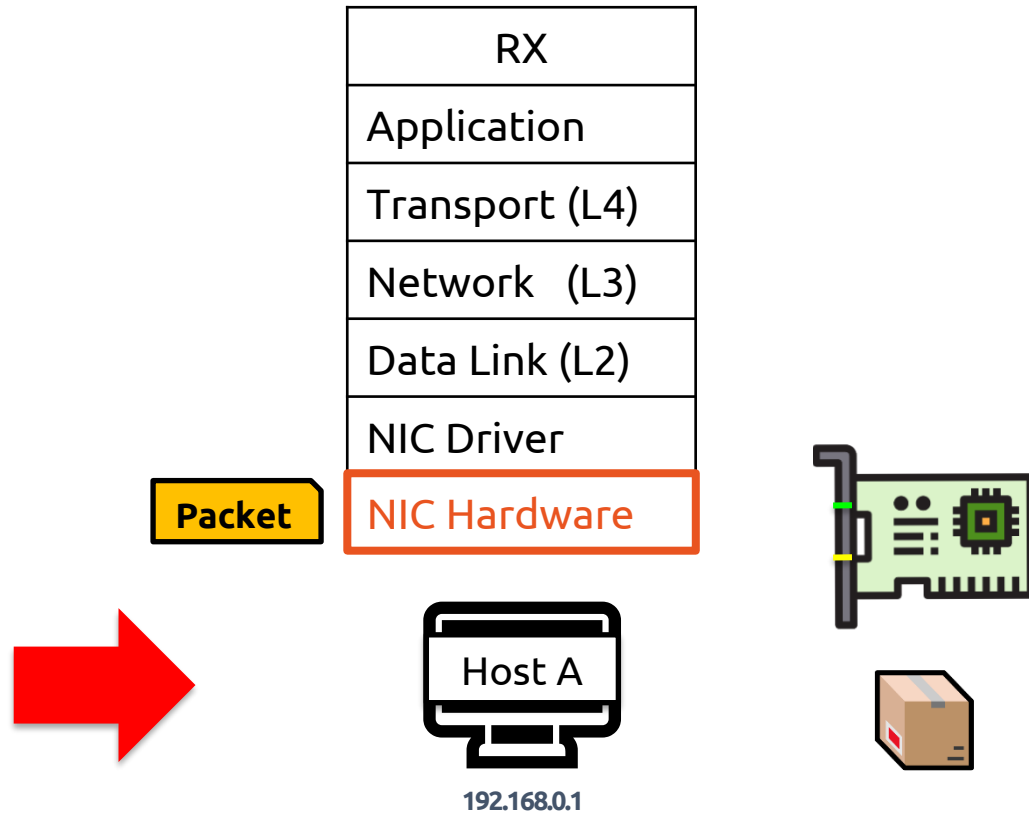
Packet Receiving Path





Packet Receiving Path

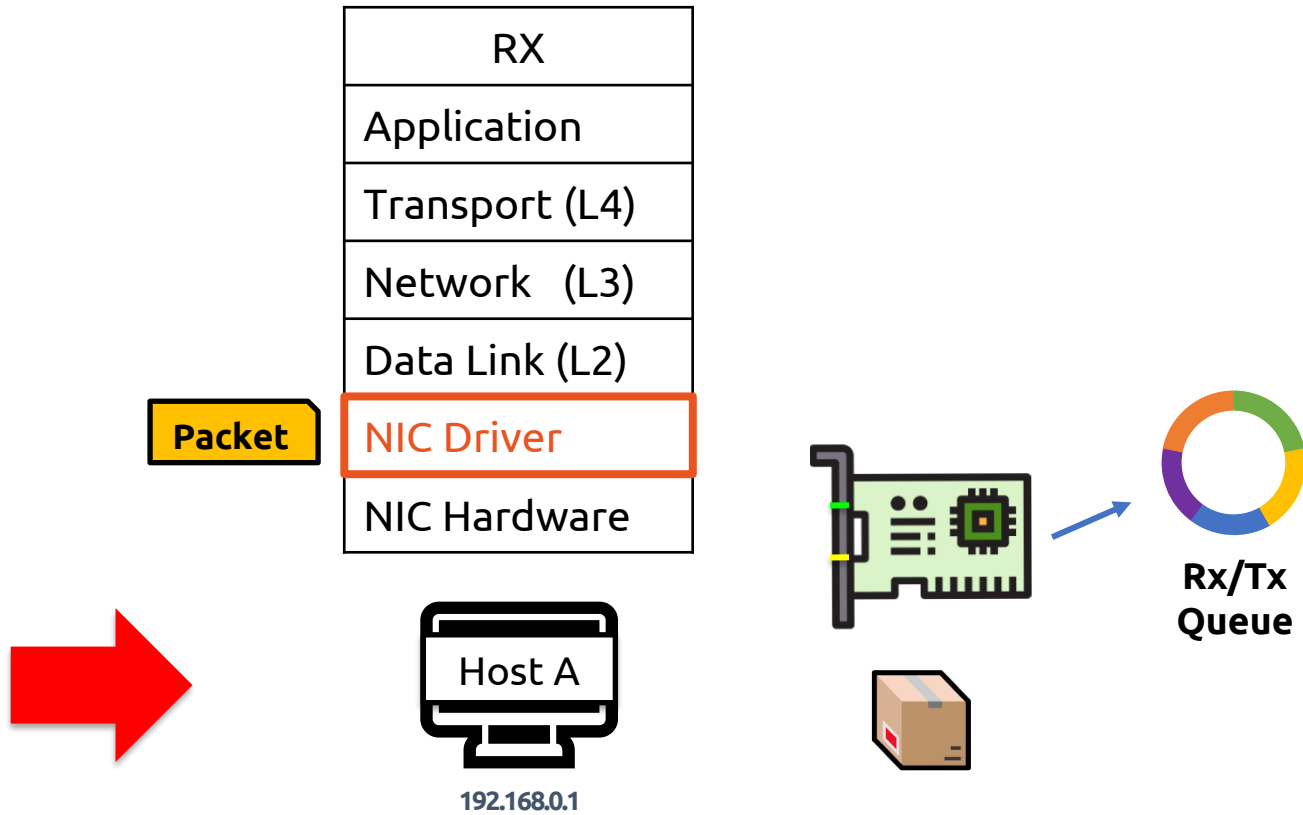
Network Interface Card receives the packet





Packet Receiving Path

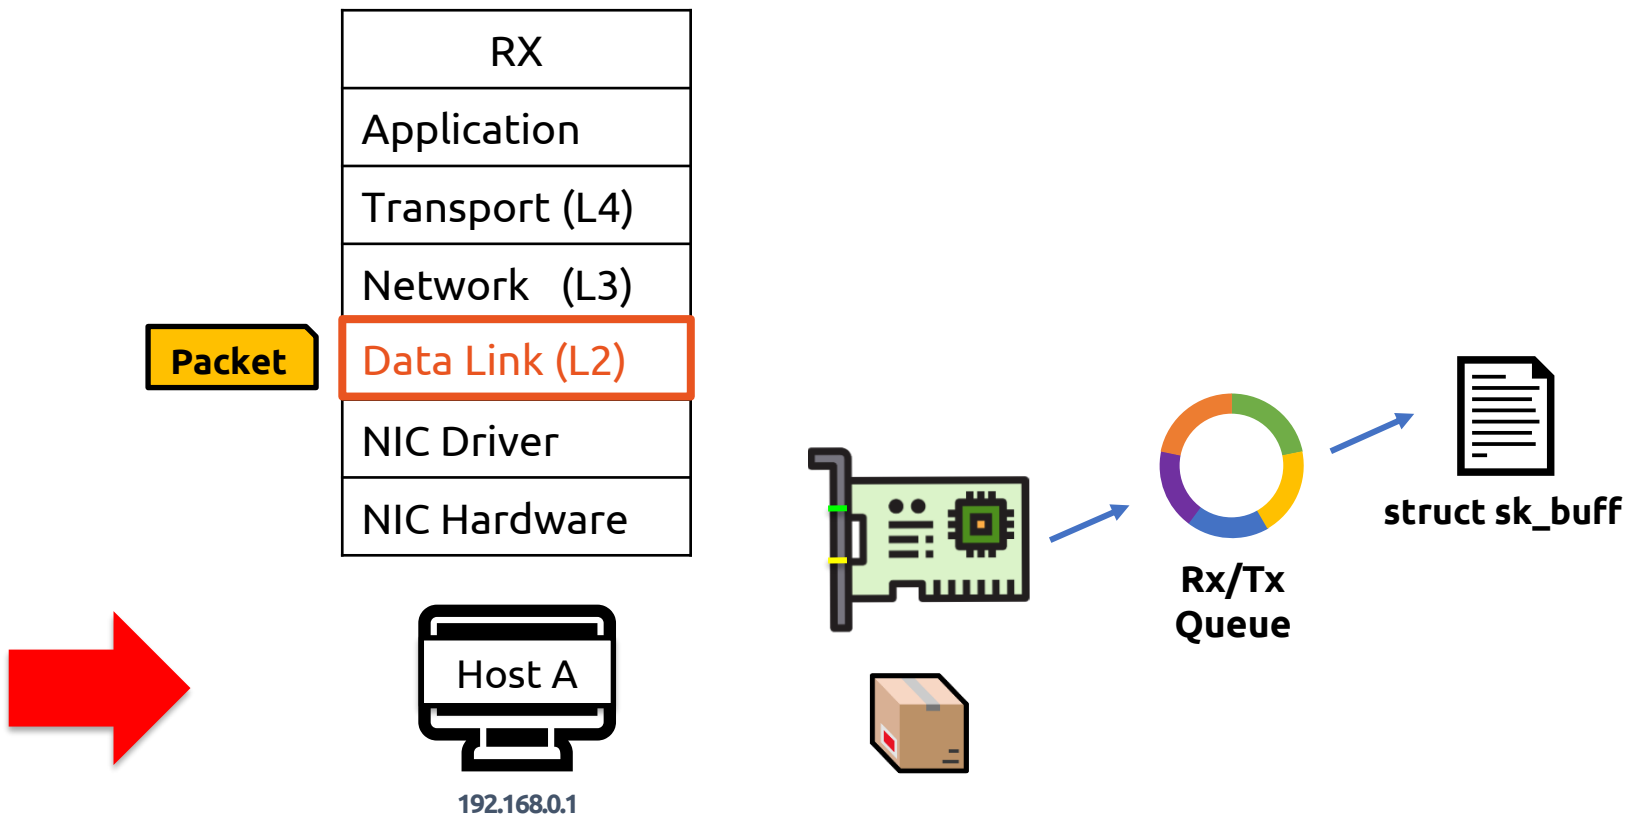
From Device Driver, pops the packet from the NIC hardware Queue (RX)





Packet Receiving Path

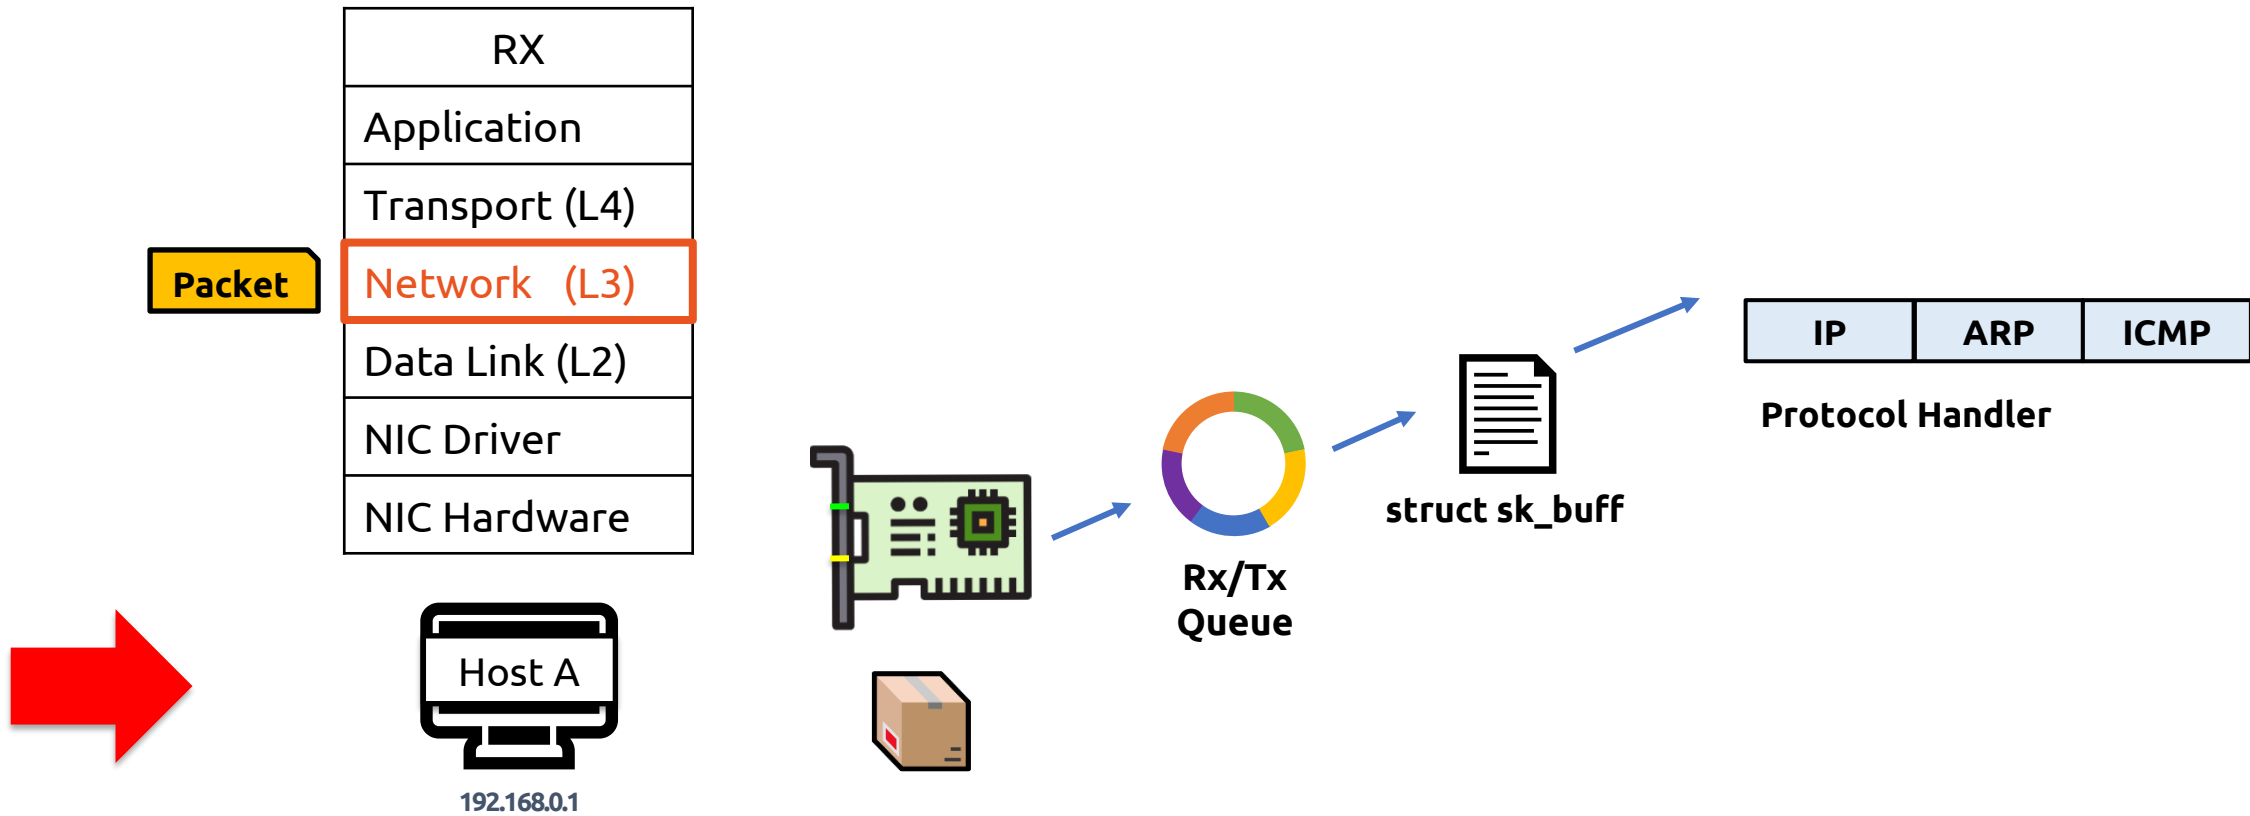
Create a `sk_buff` structure with this packet.





Packet Receiving Path

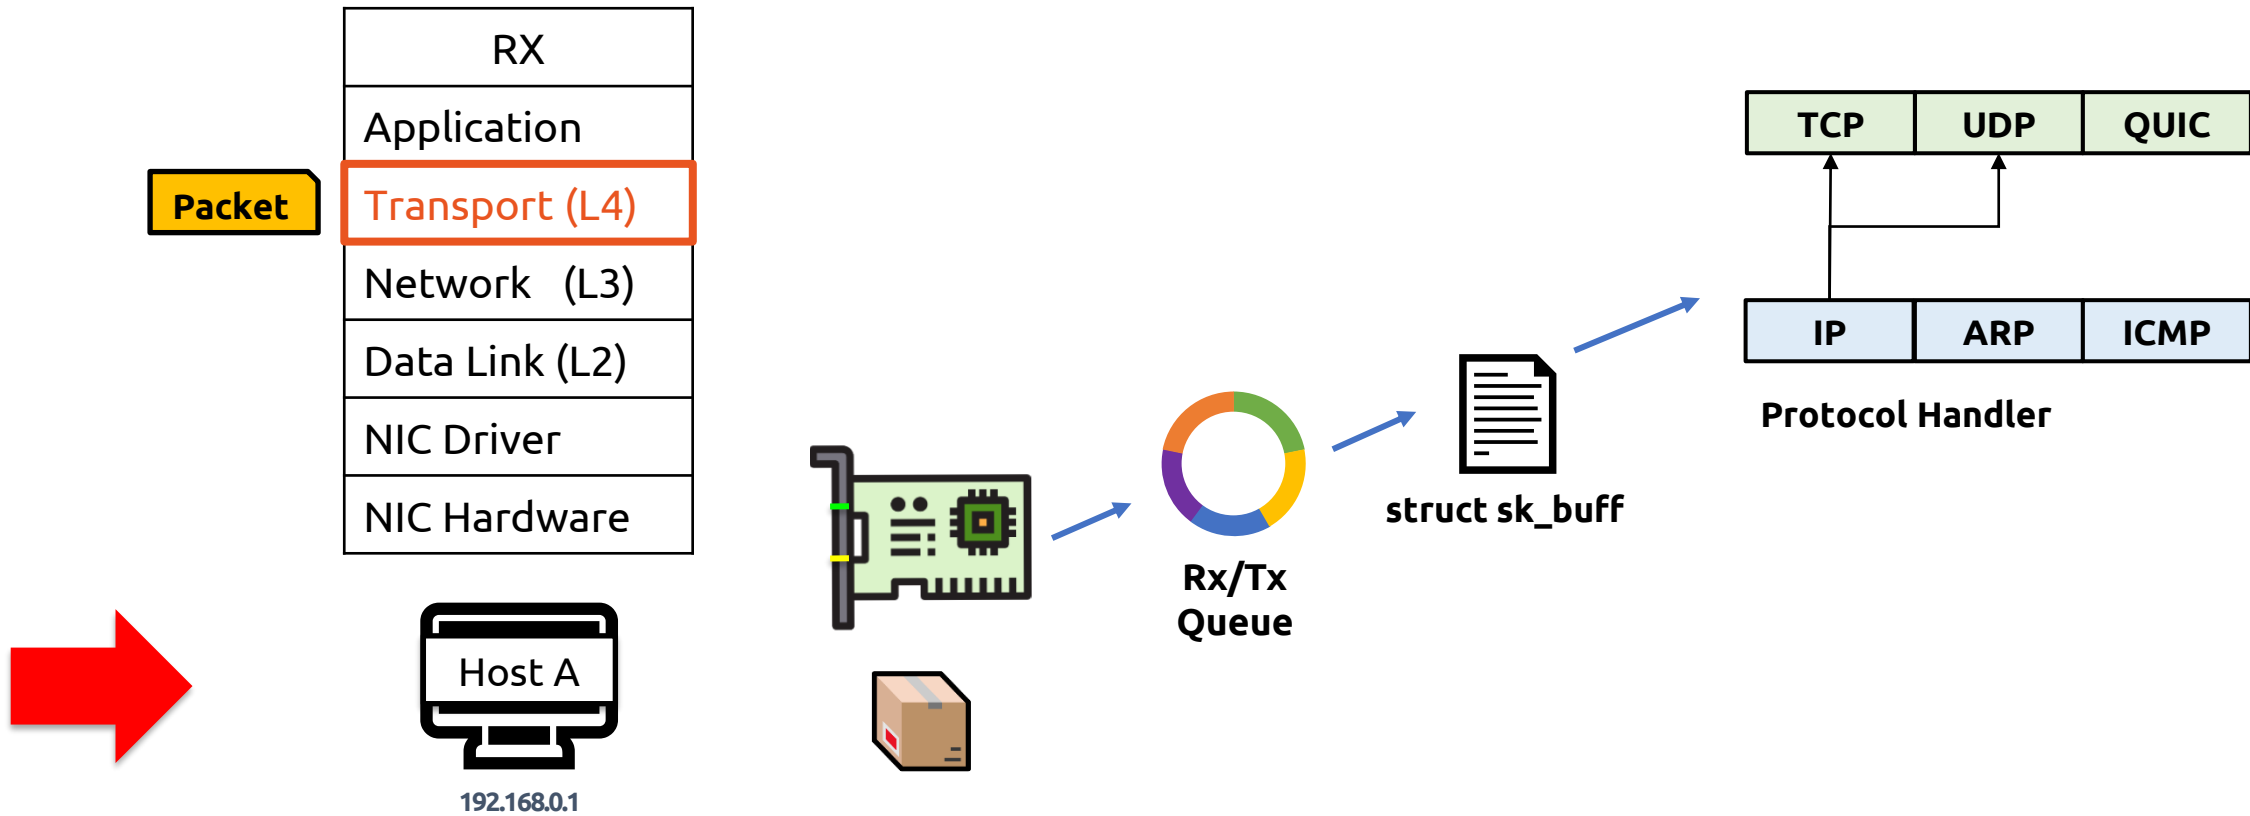
Determines which protocol the packet is,





Packet Receiving Path

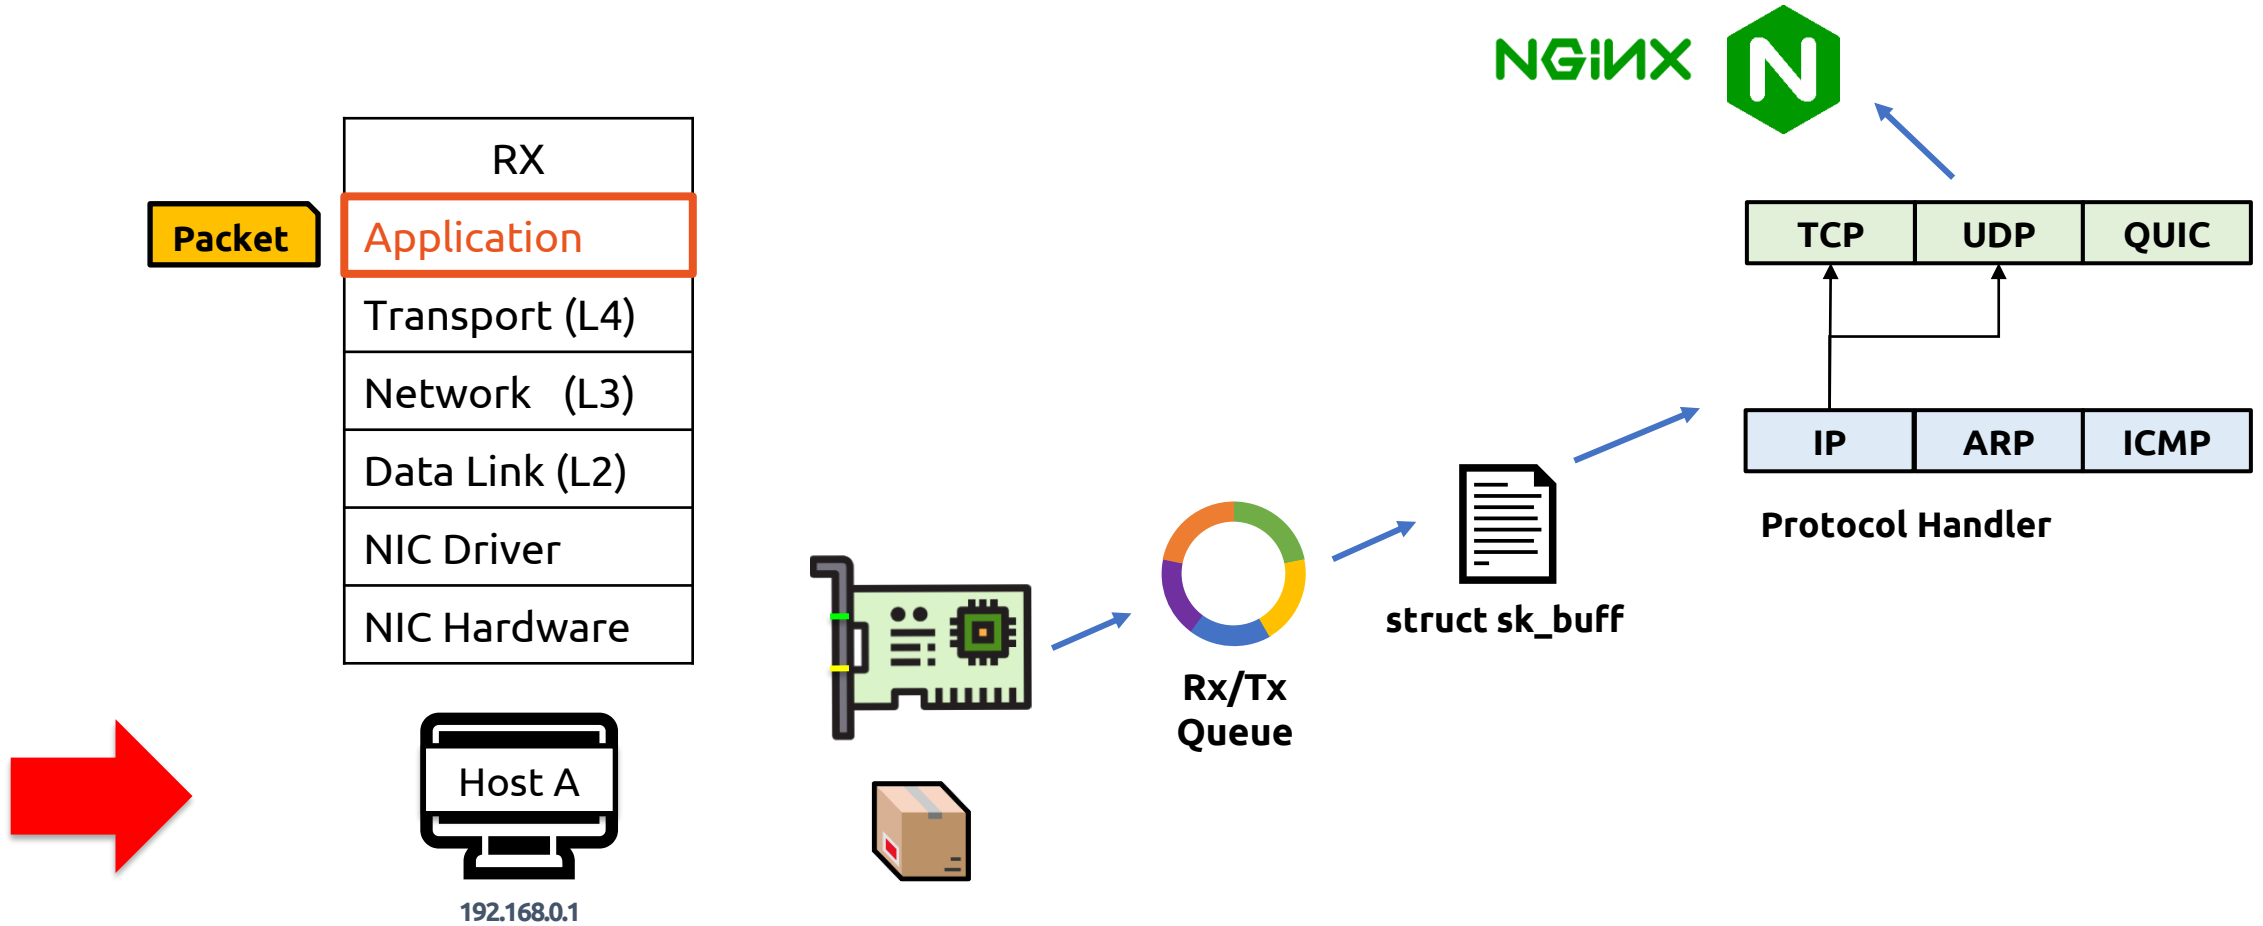
Call the appropriate handler function for the protocol





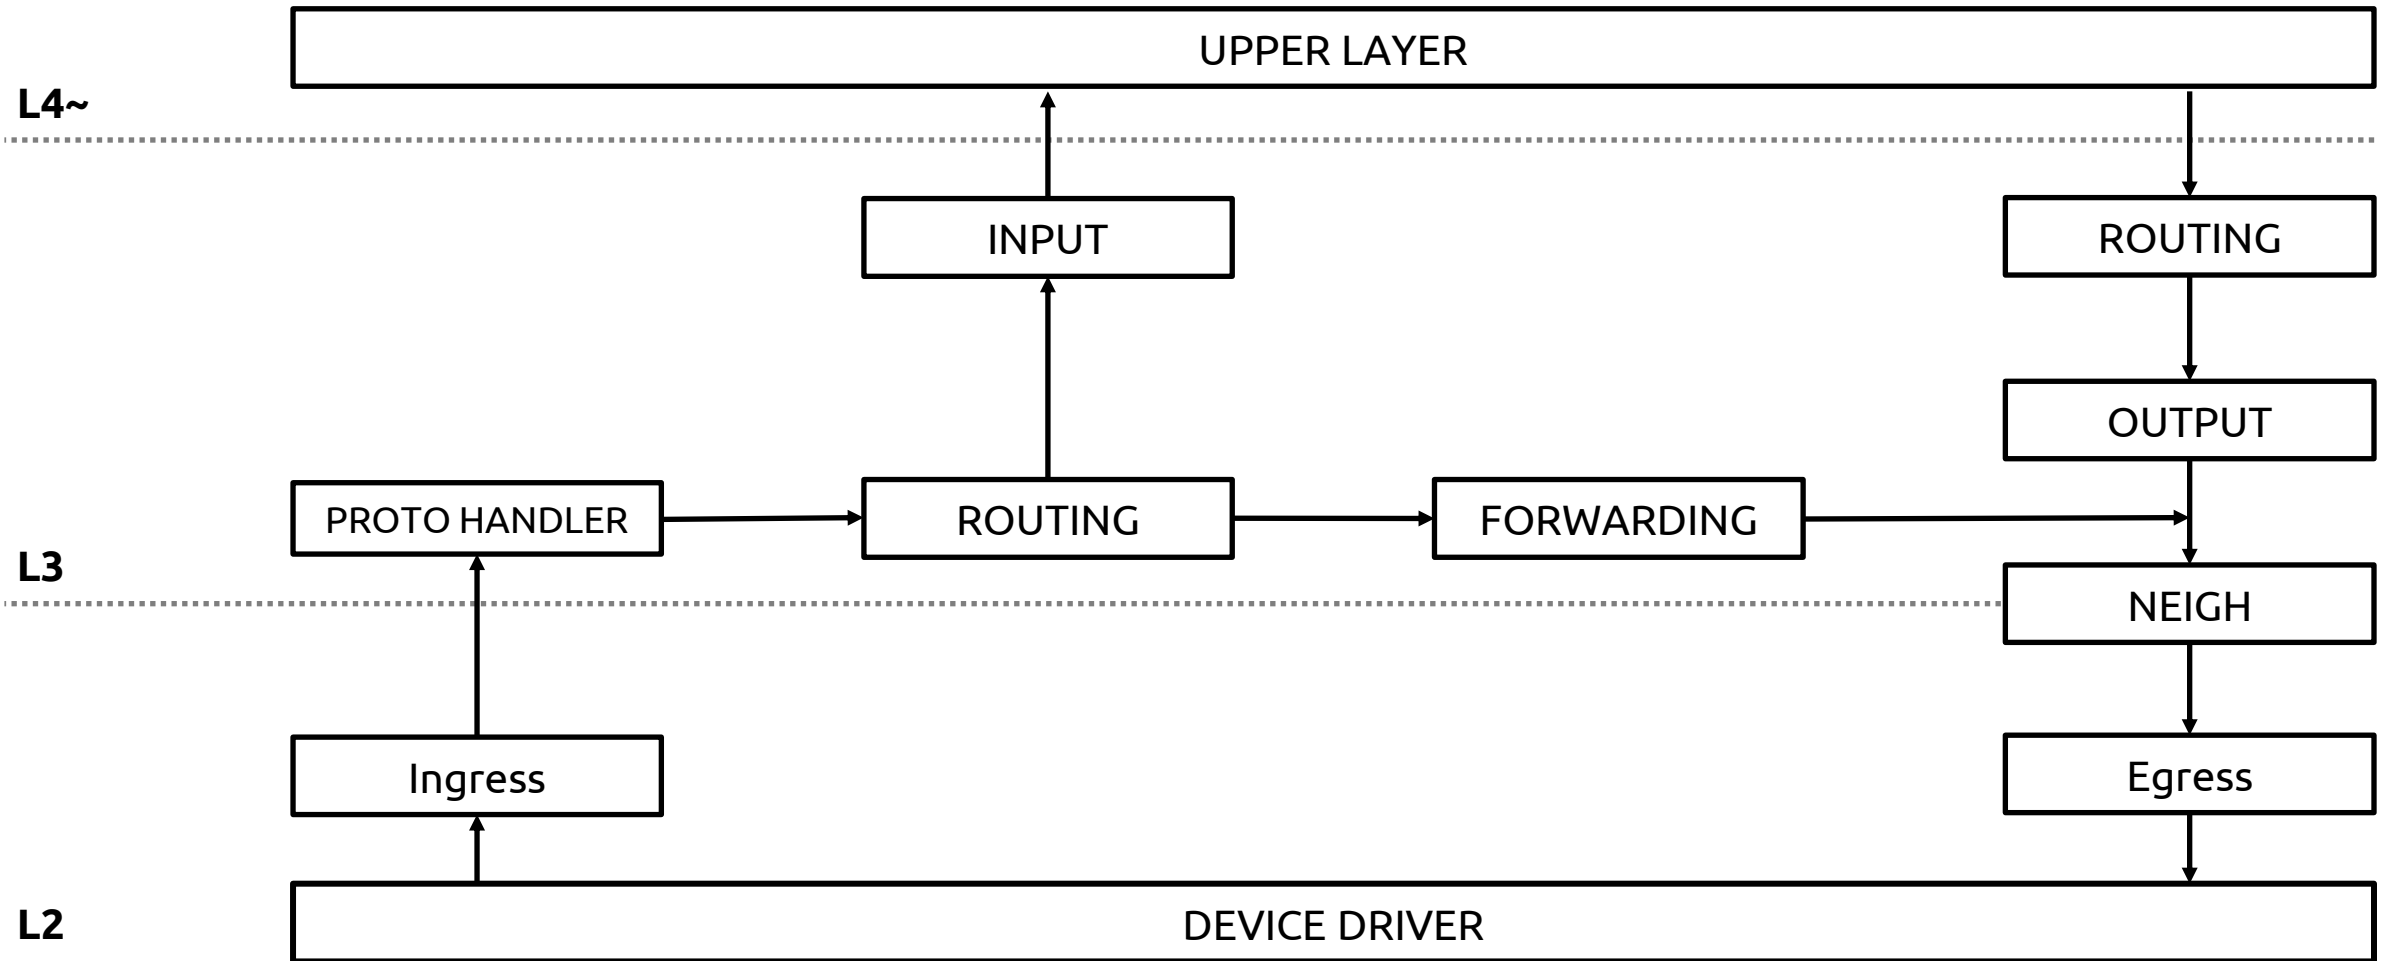
Packet Receiving Path

Hand over processed packet to user applications






Linux Kernel Networking Packet Path





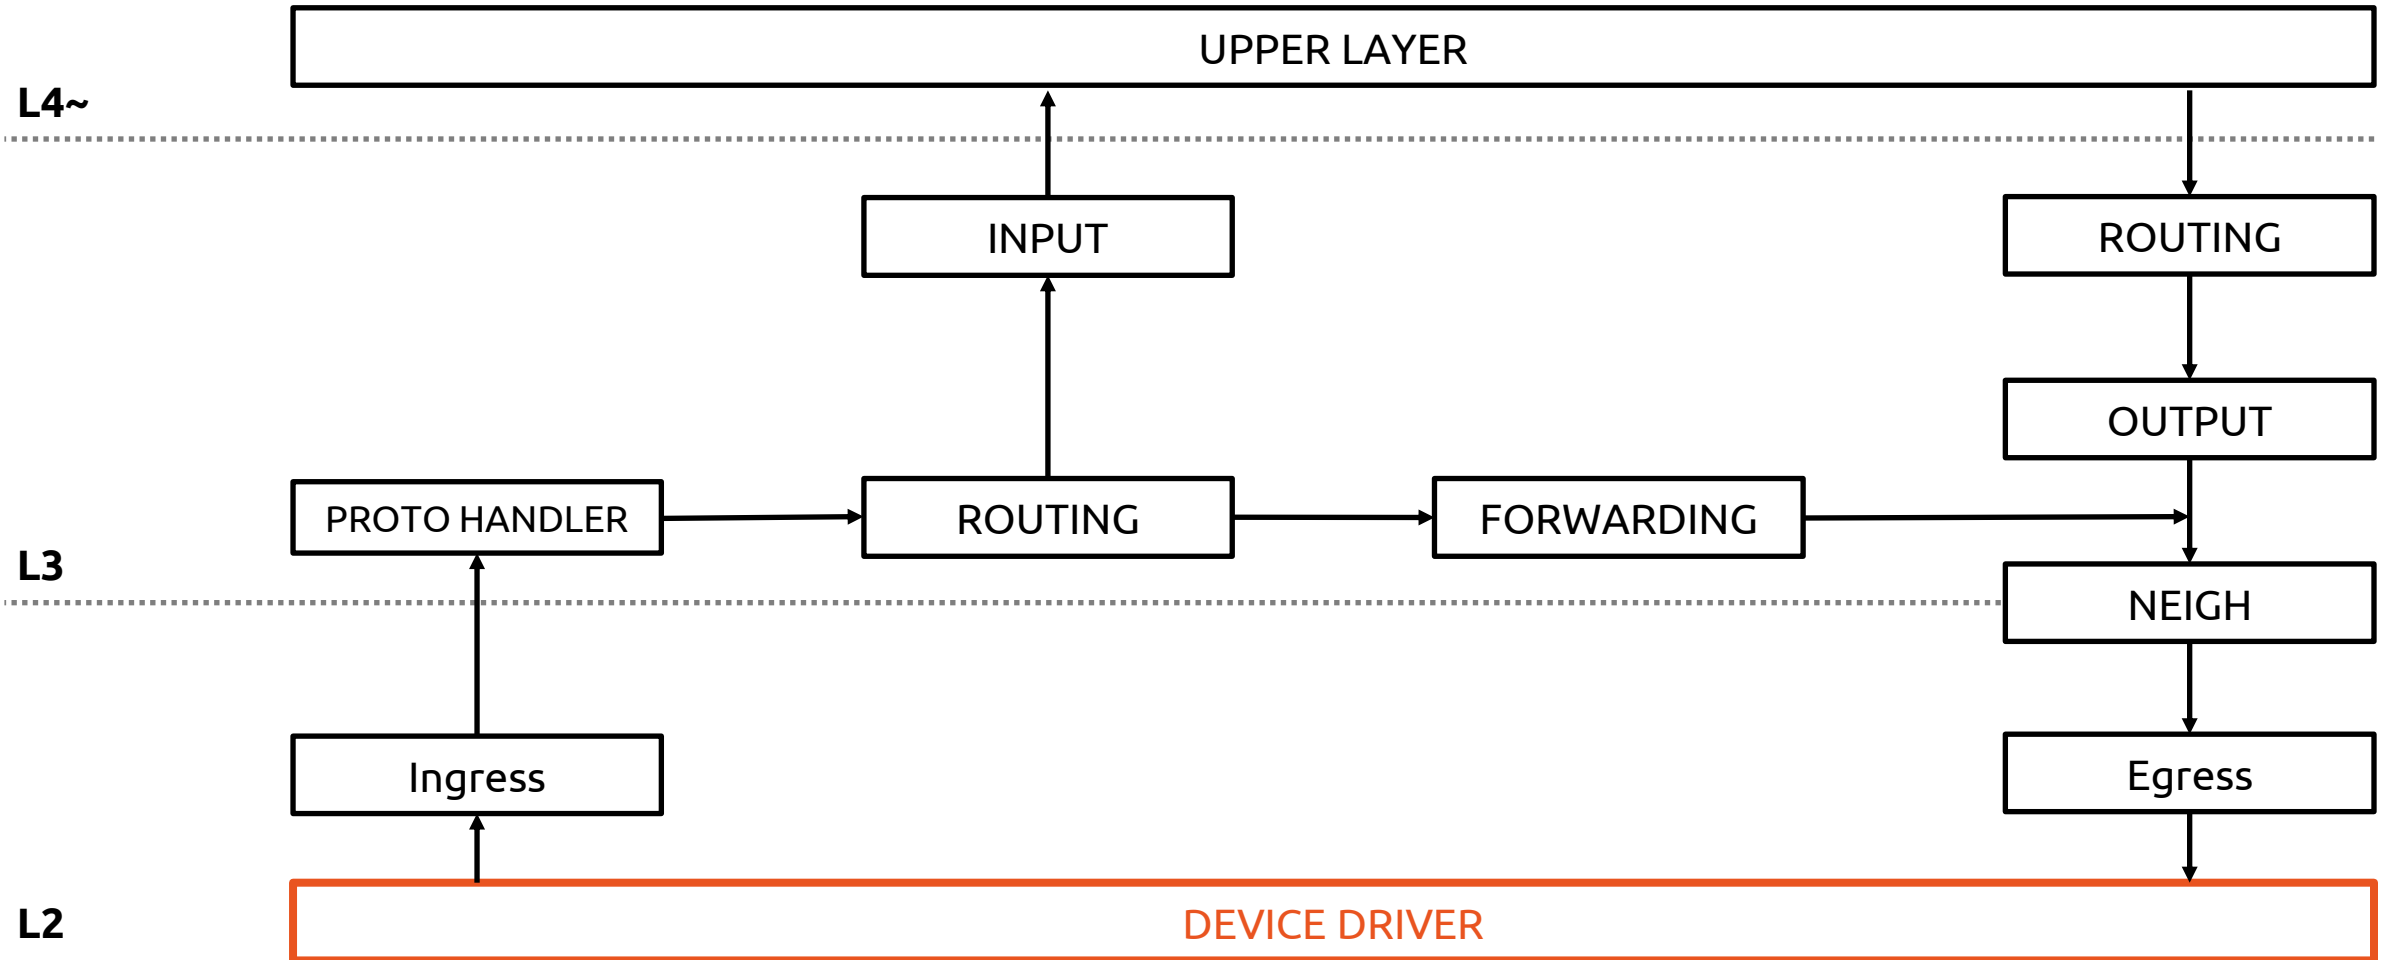
Packet Receiving Path





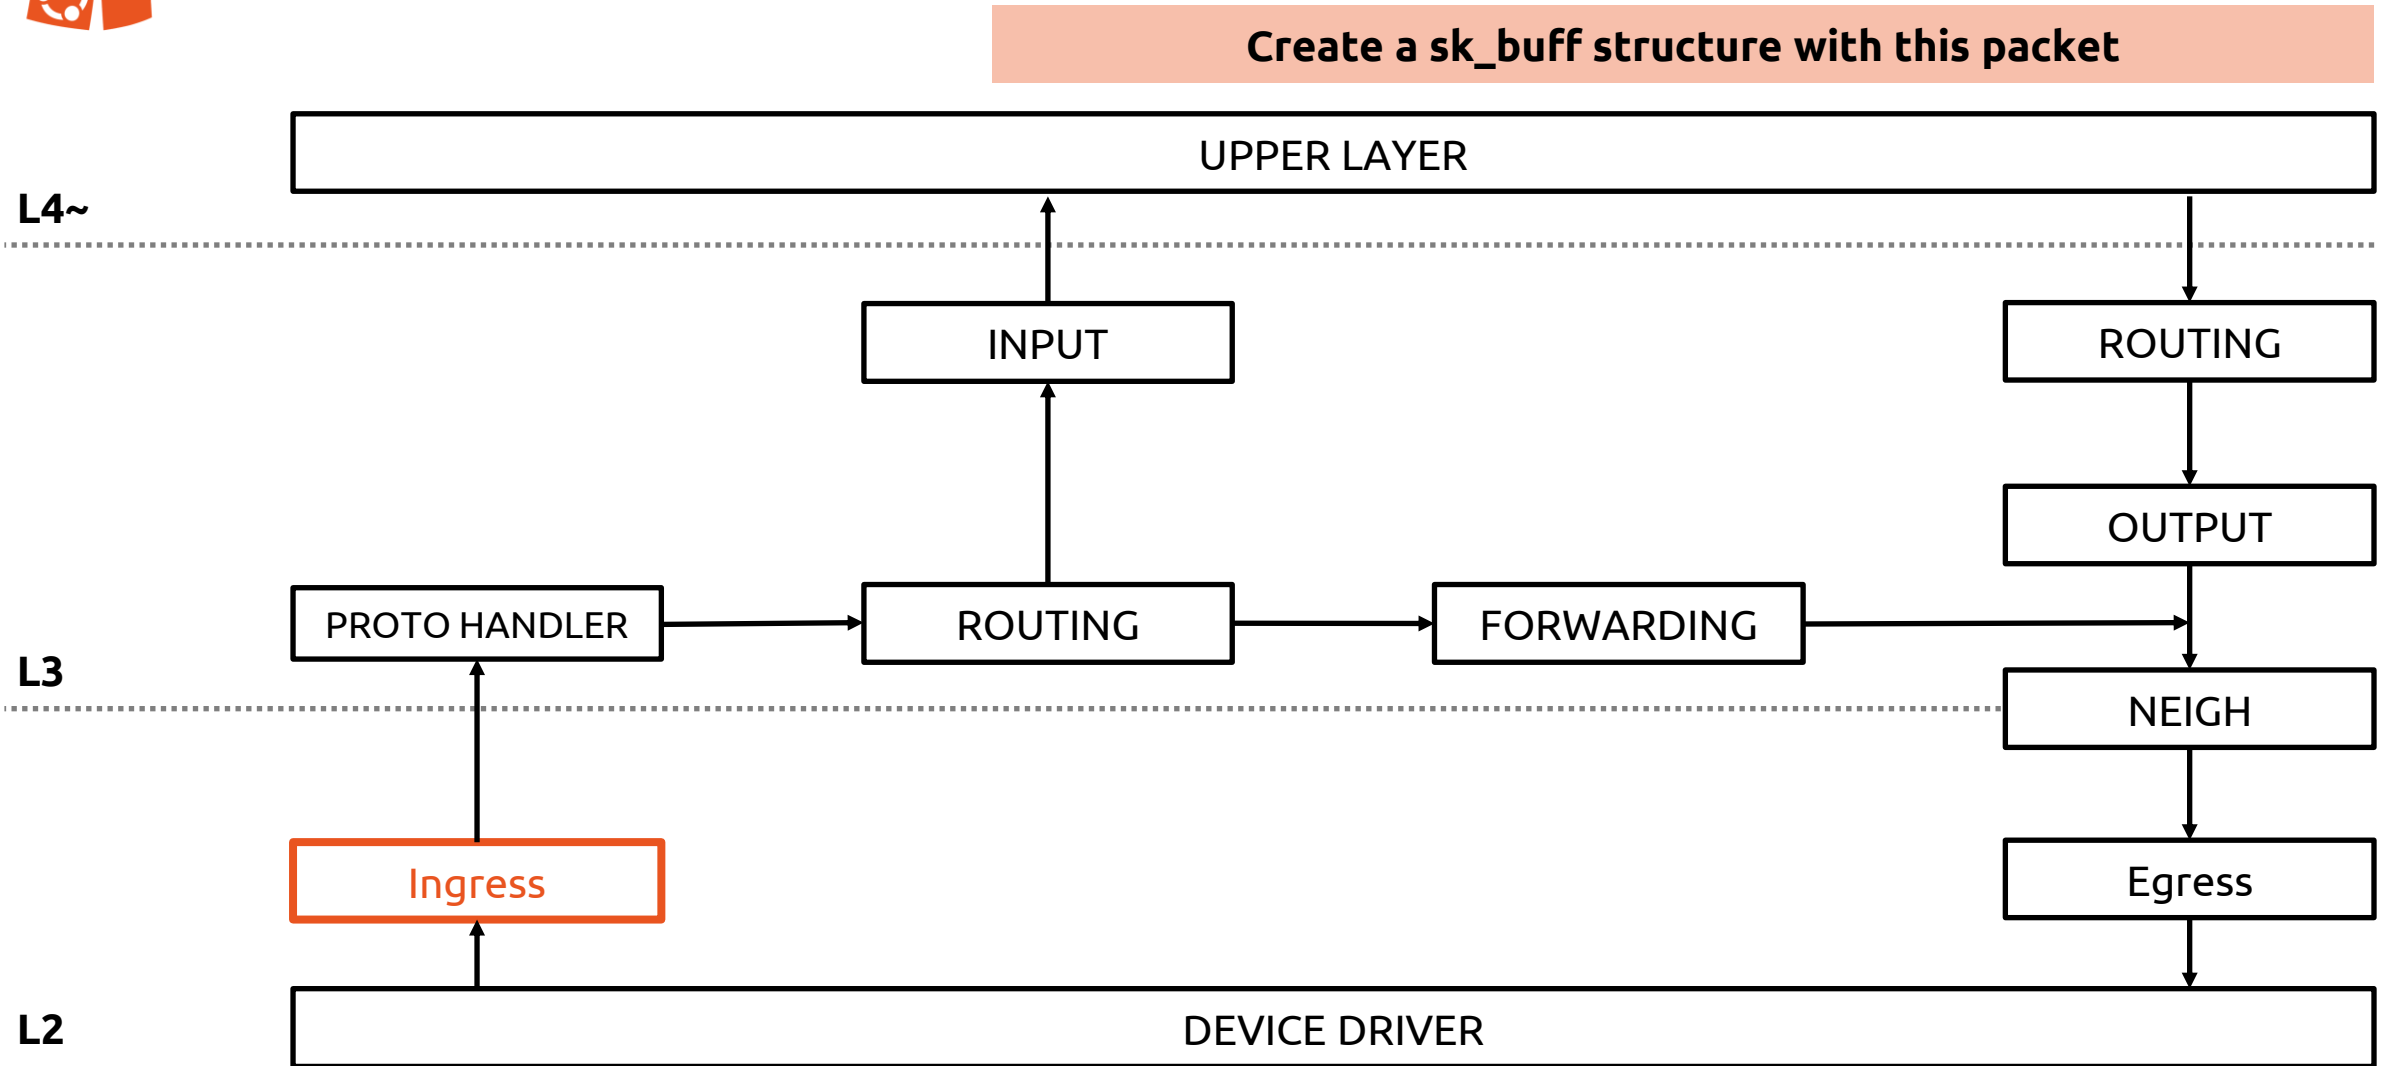
Packet Receiving Path

After Network Interface Card receives the packet, Device Driver pops the packet from the NIC hardware Queue (RX)





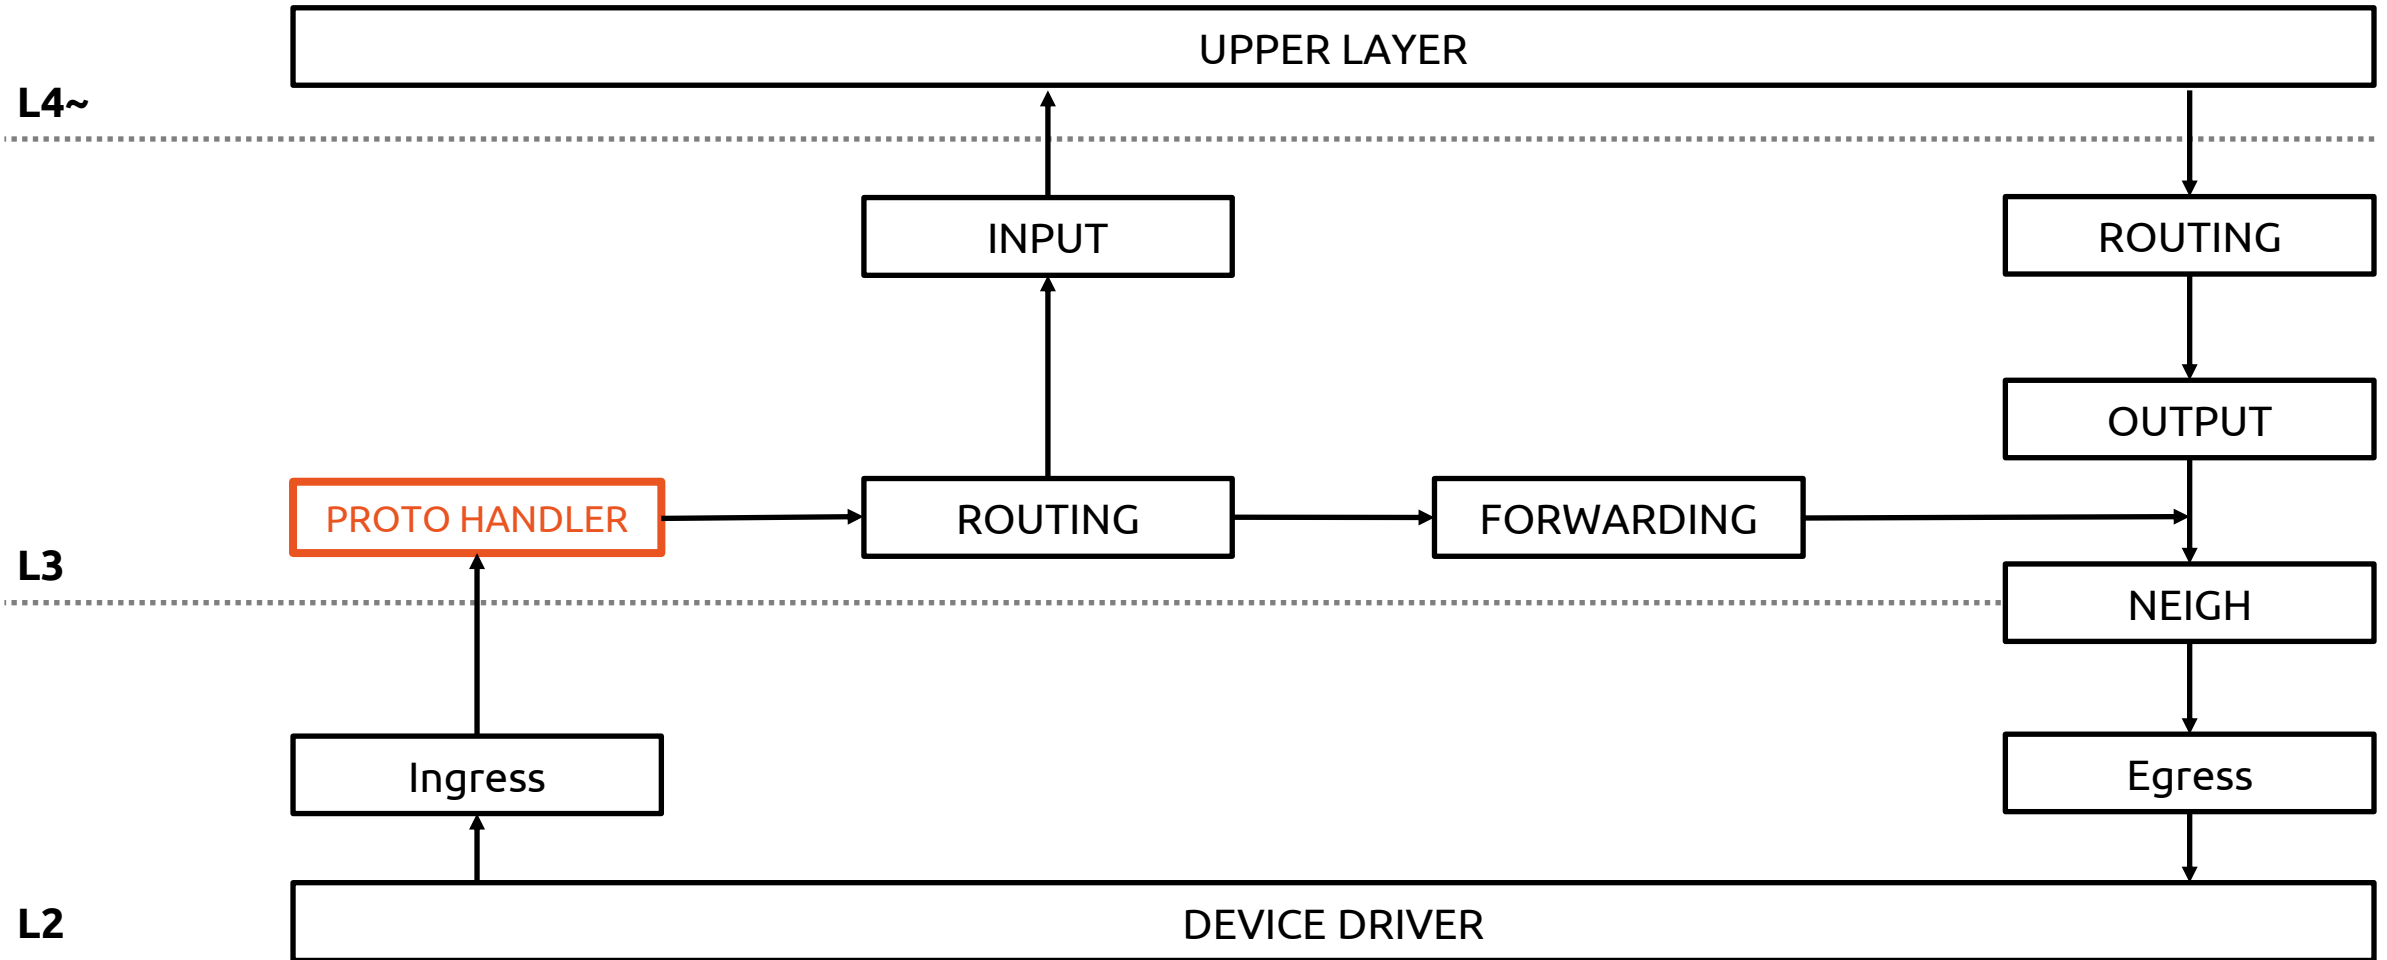
Packet Receiving Path





Packet Receiving Path


Determine which protocol the packet is and call the appropriate handler function for the protocol





Packet Receiving Path

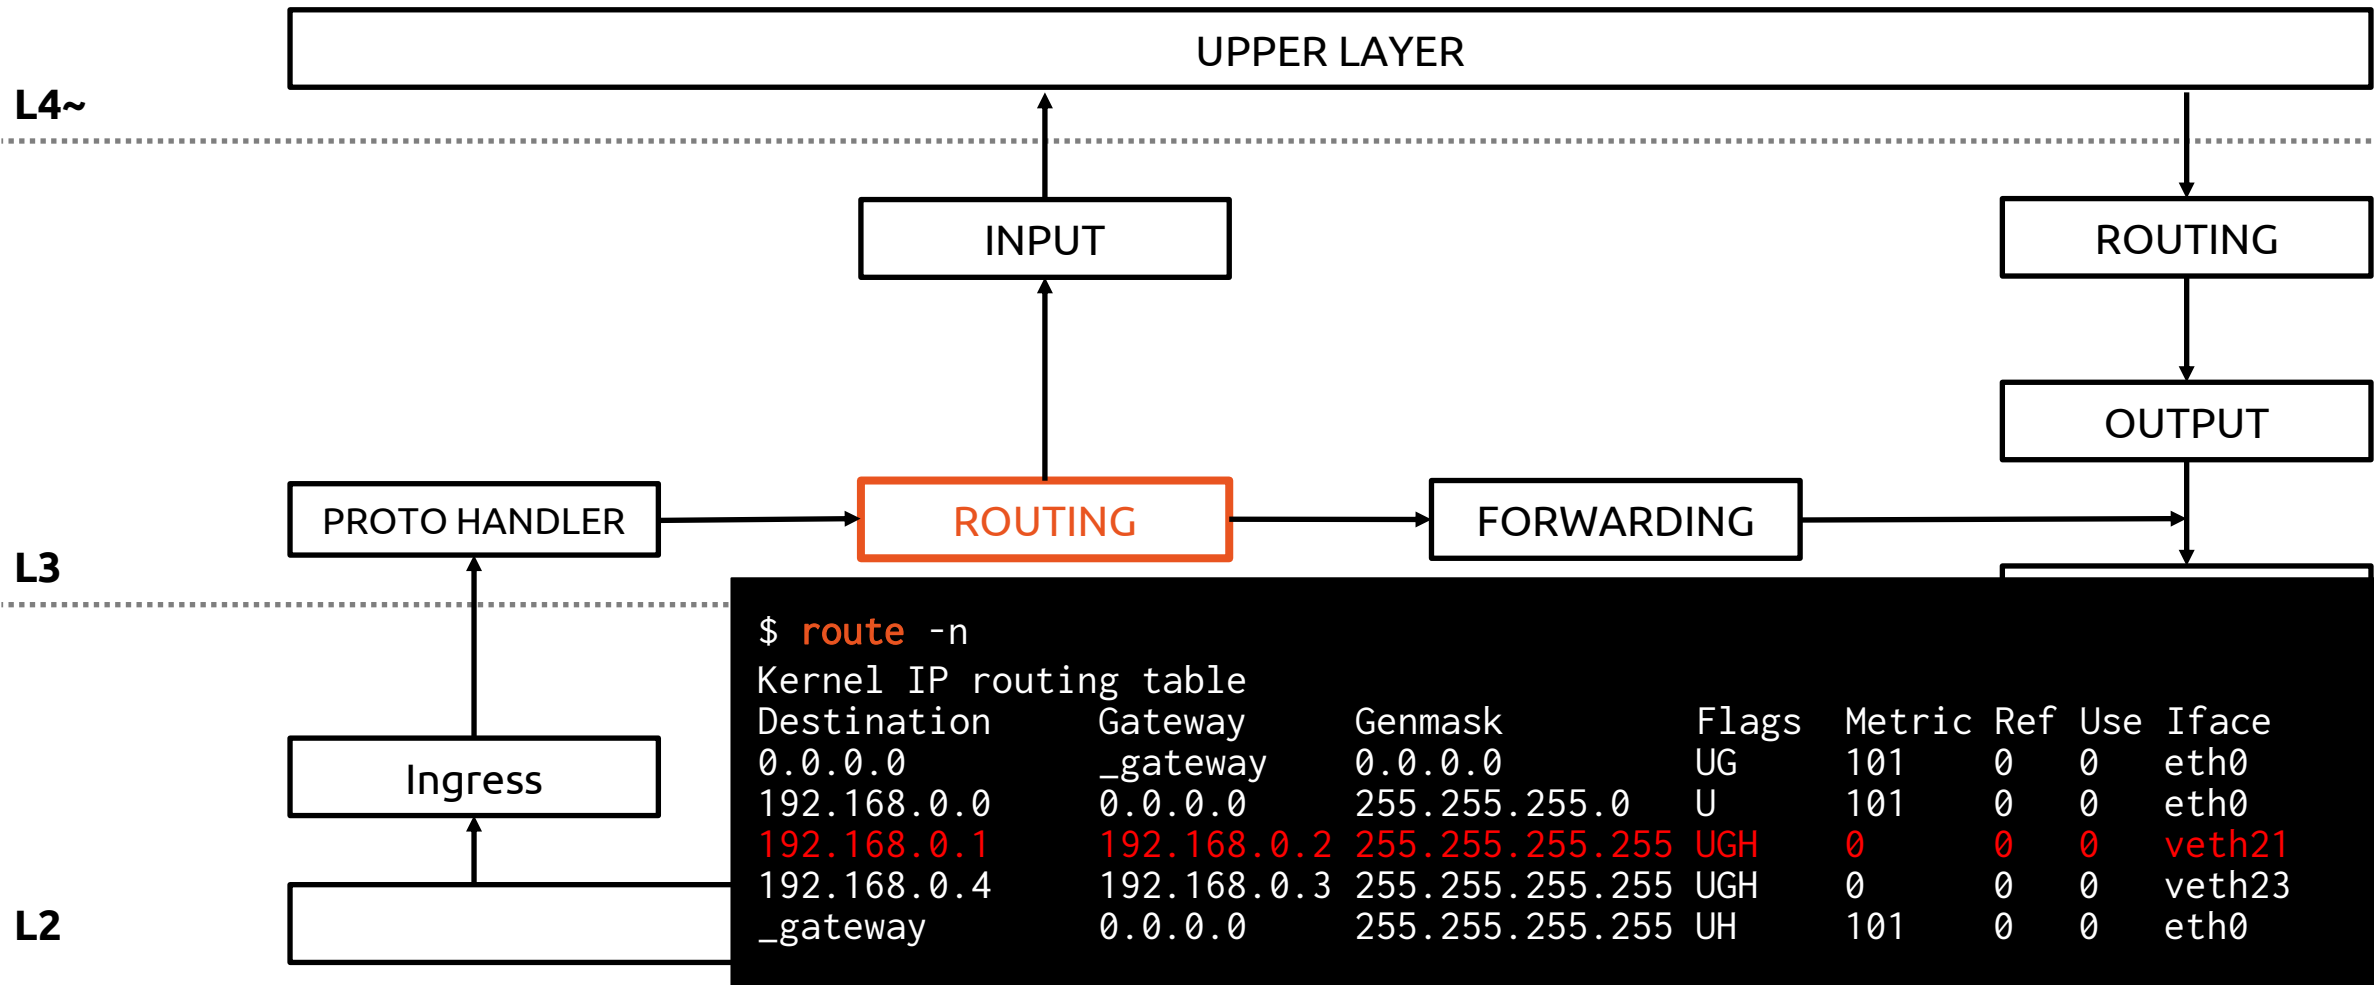
Lookup the ROUTING table to identify where the packet is headed





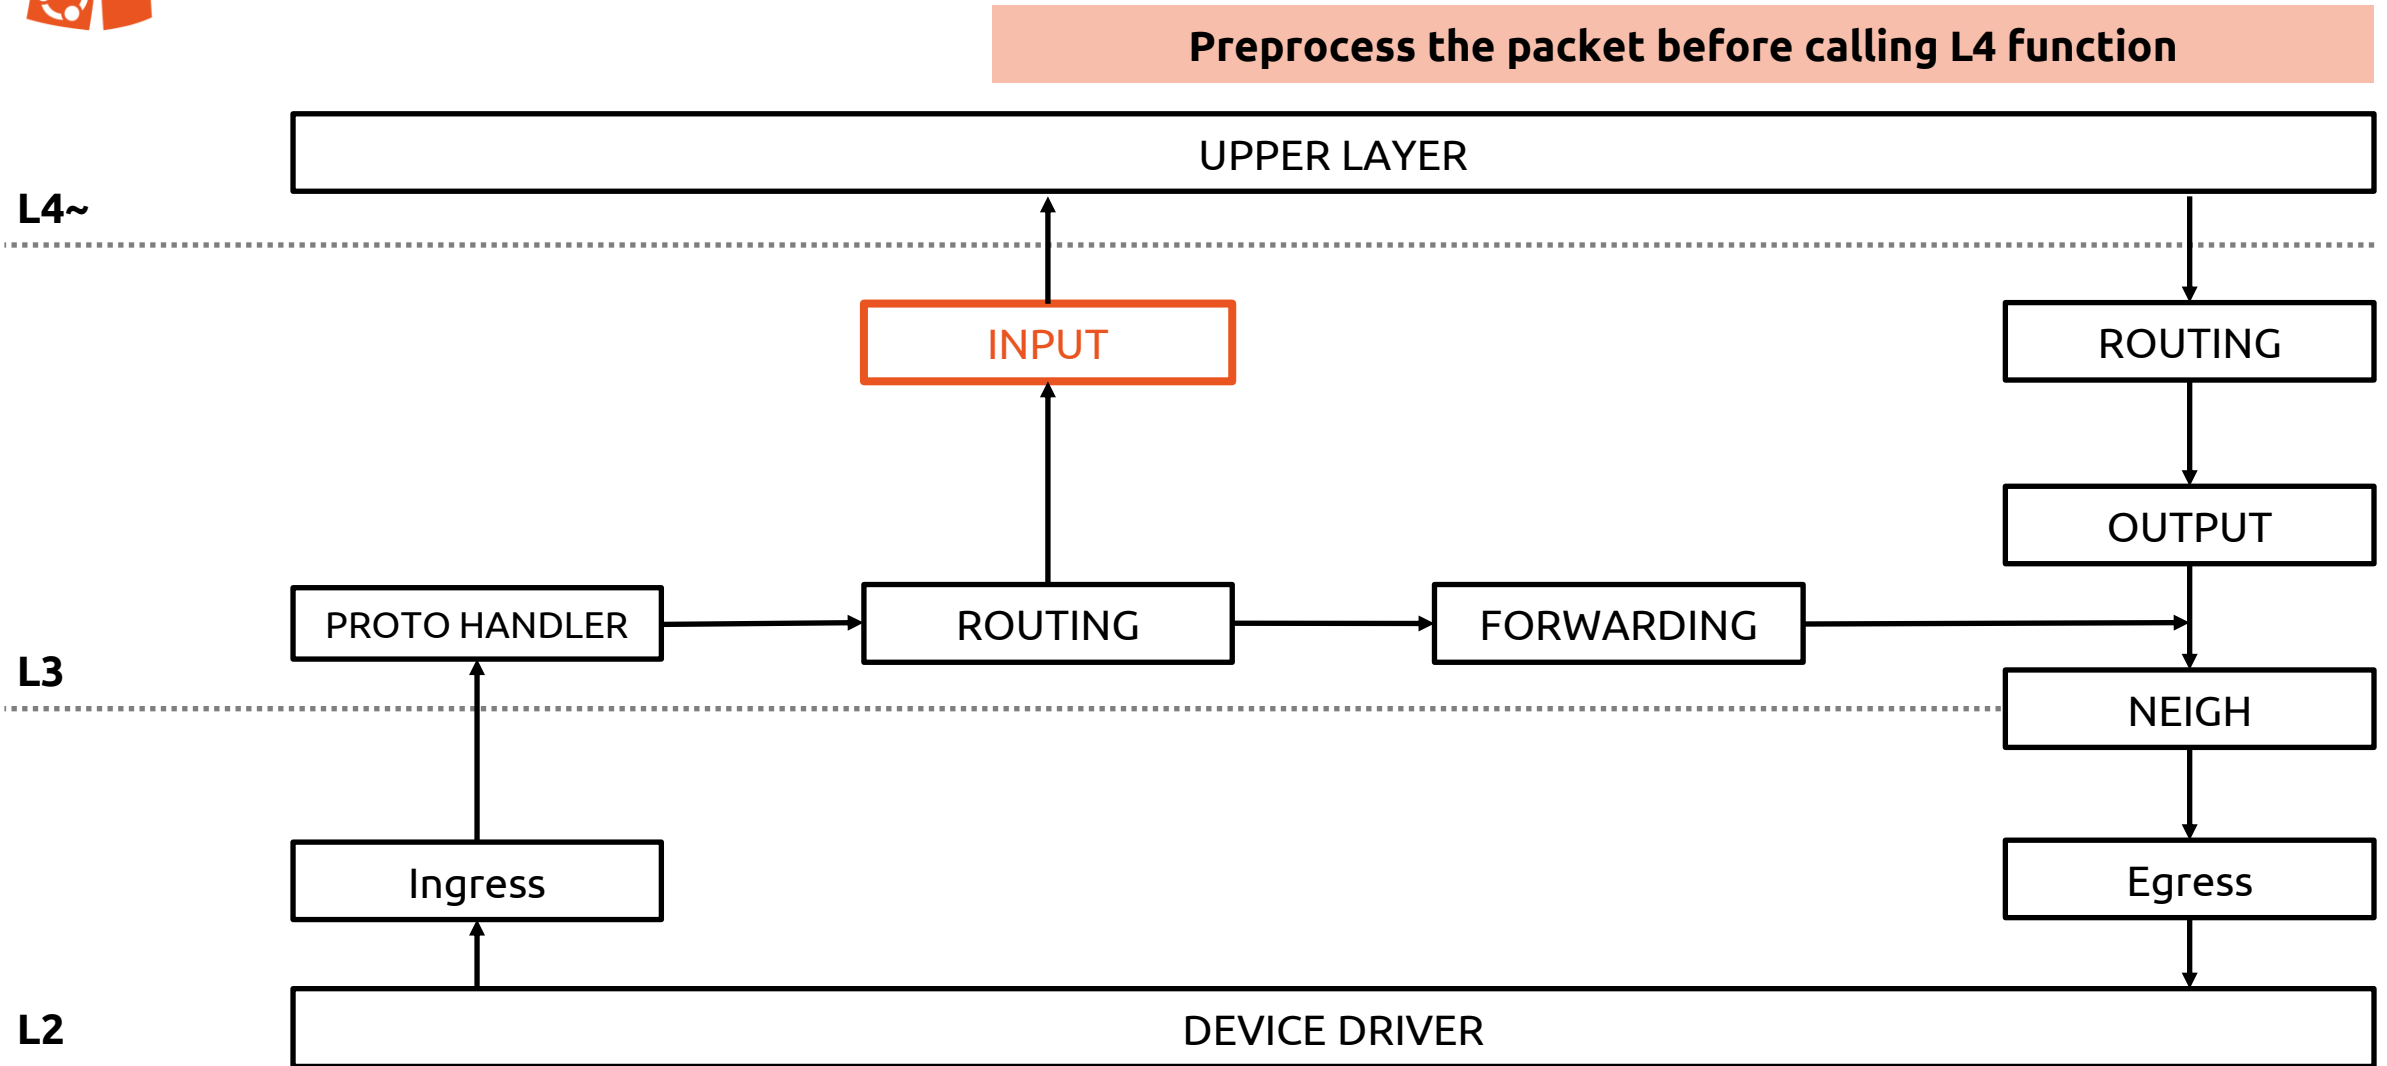
Packet Receiving Path

Lookup the ROUTING table to identify where the packet is headed



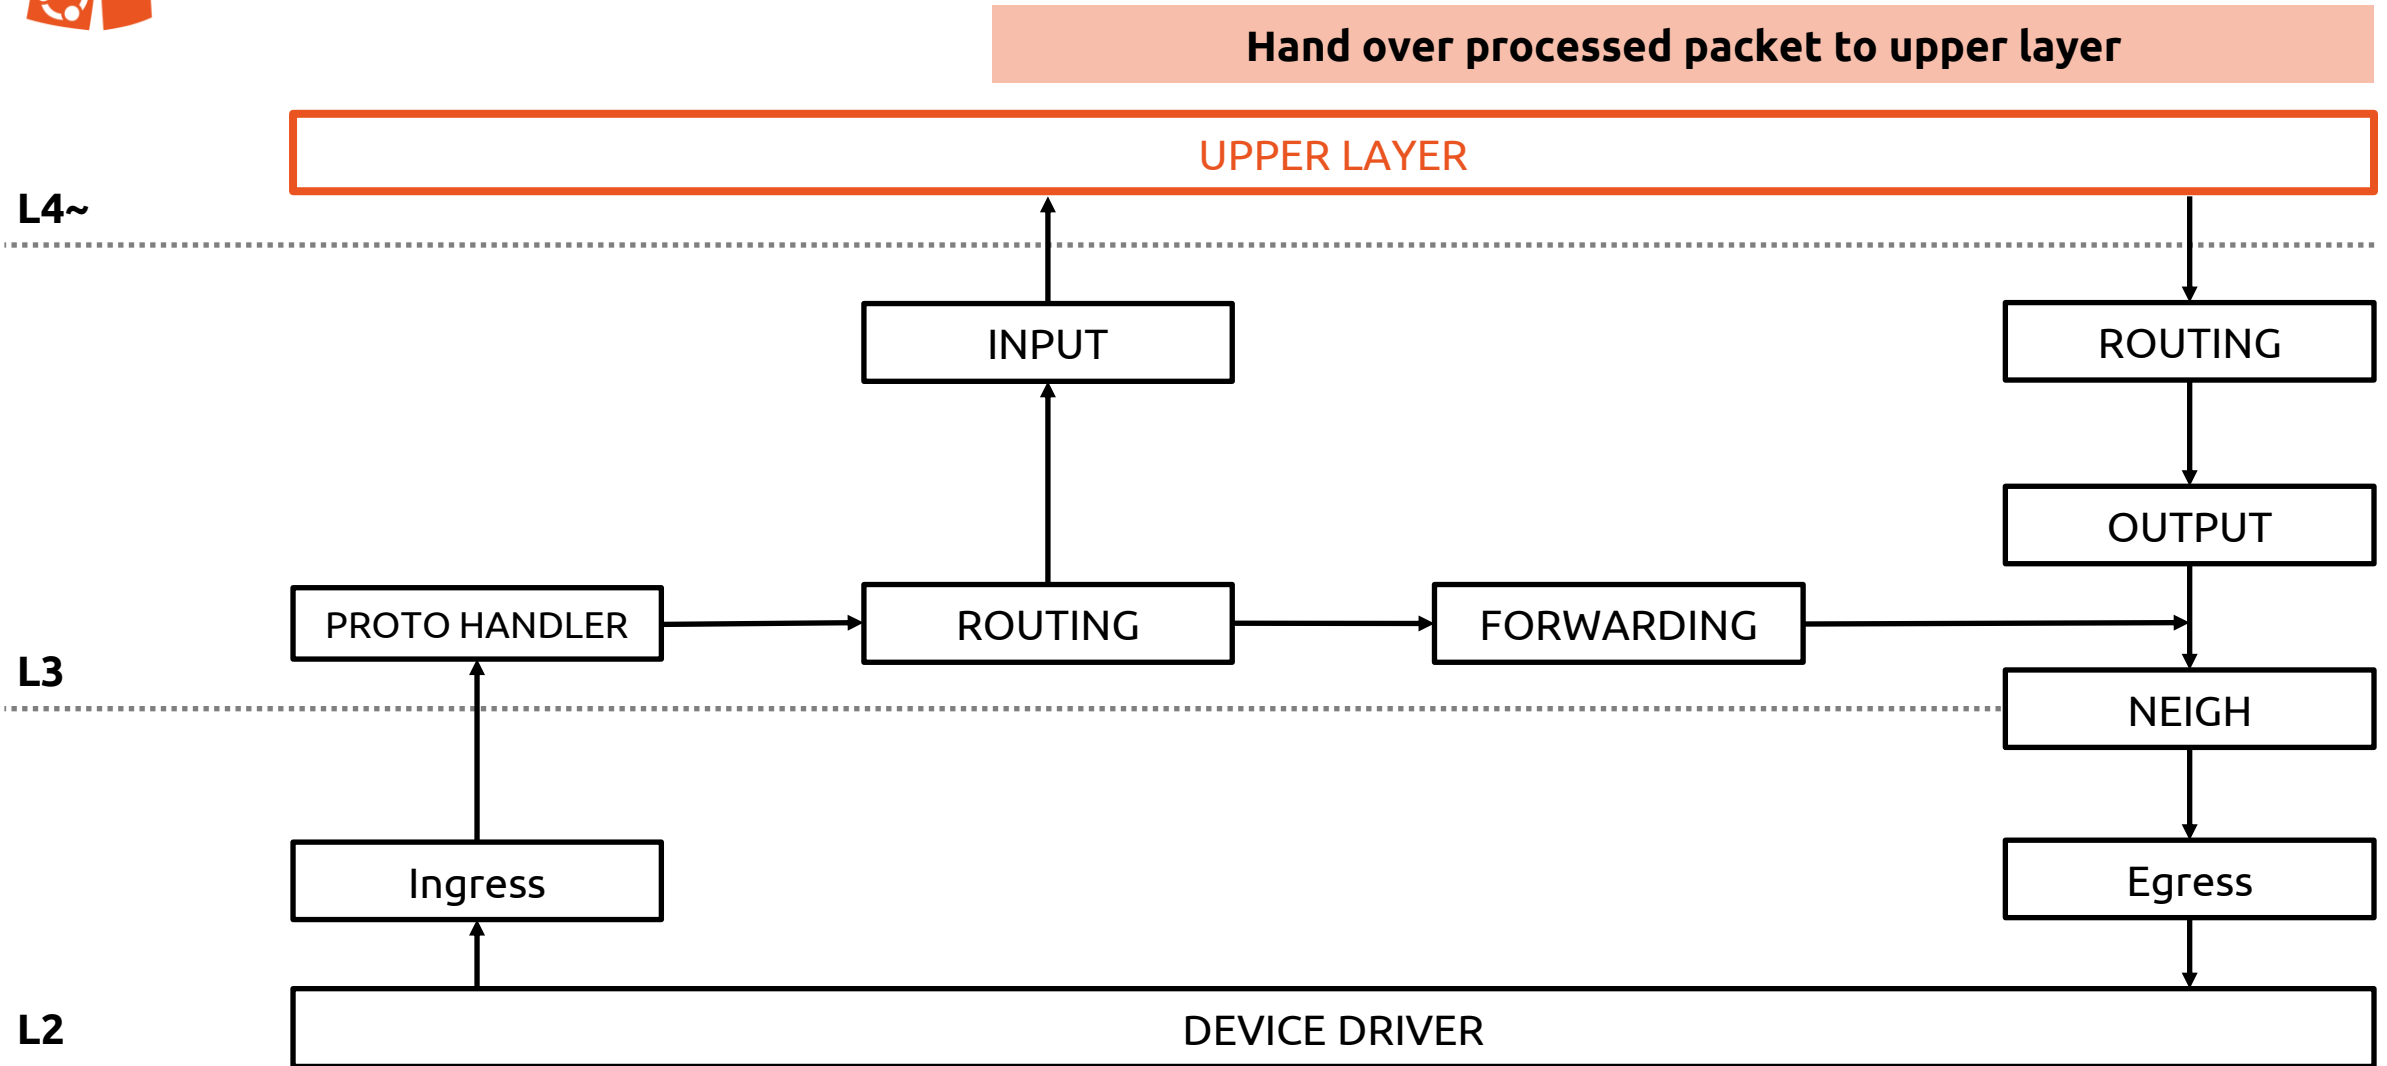


Packet Receiving Path



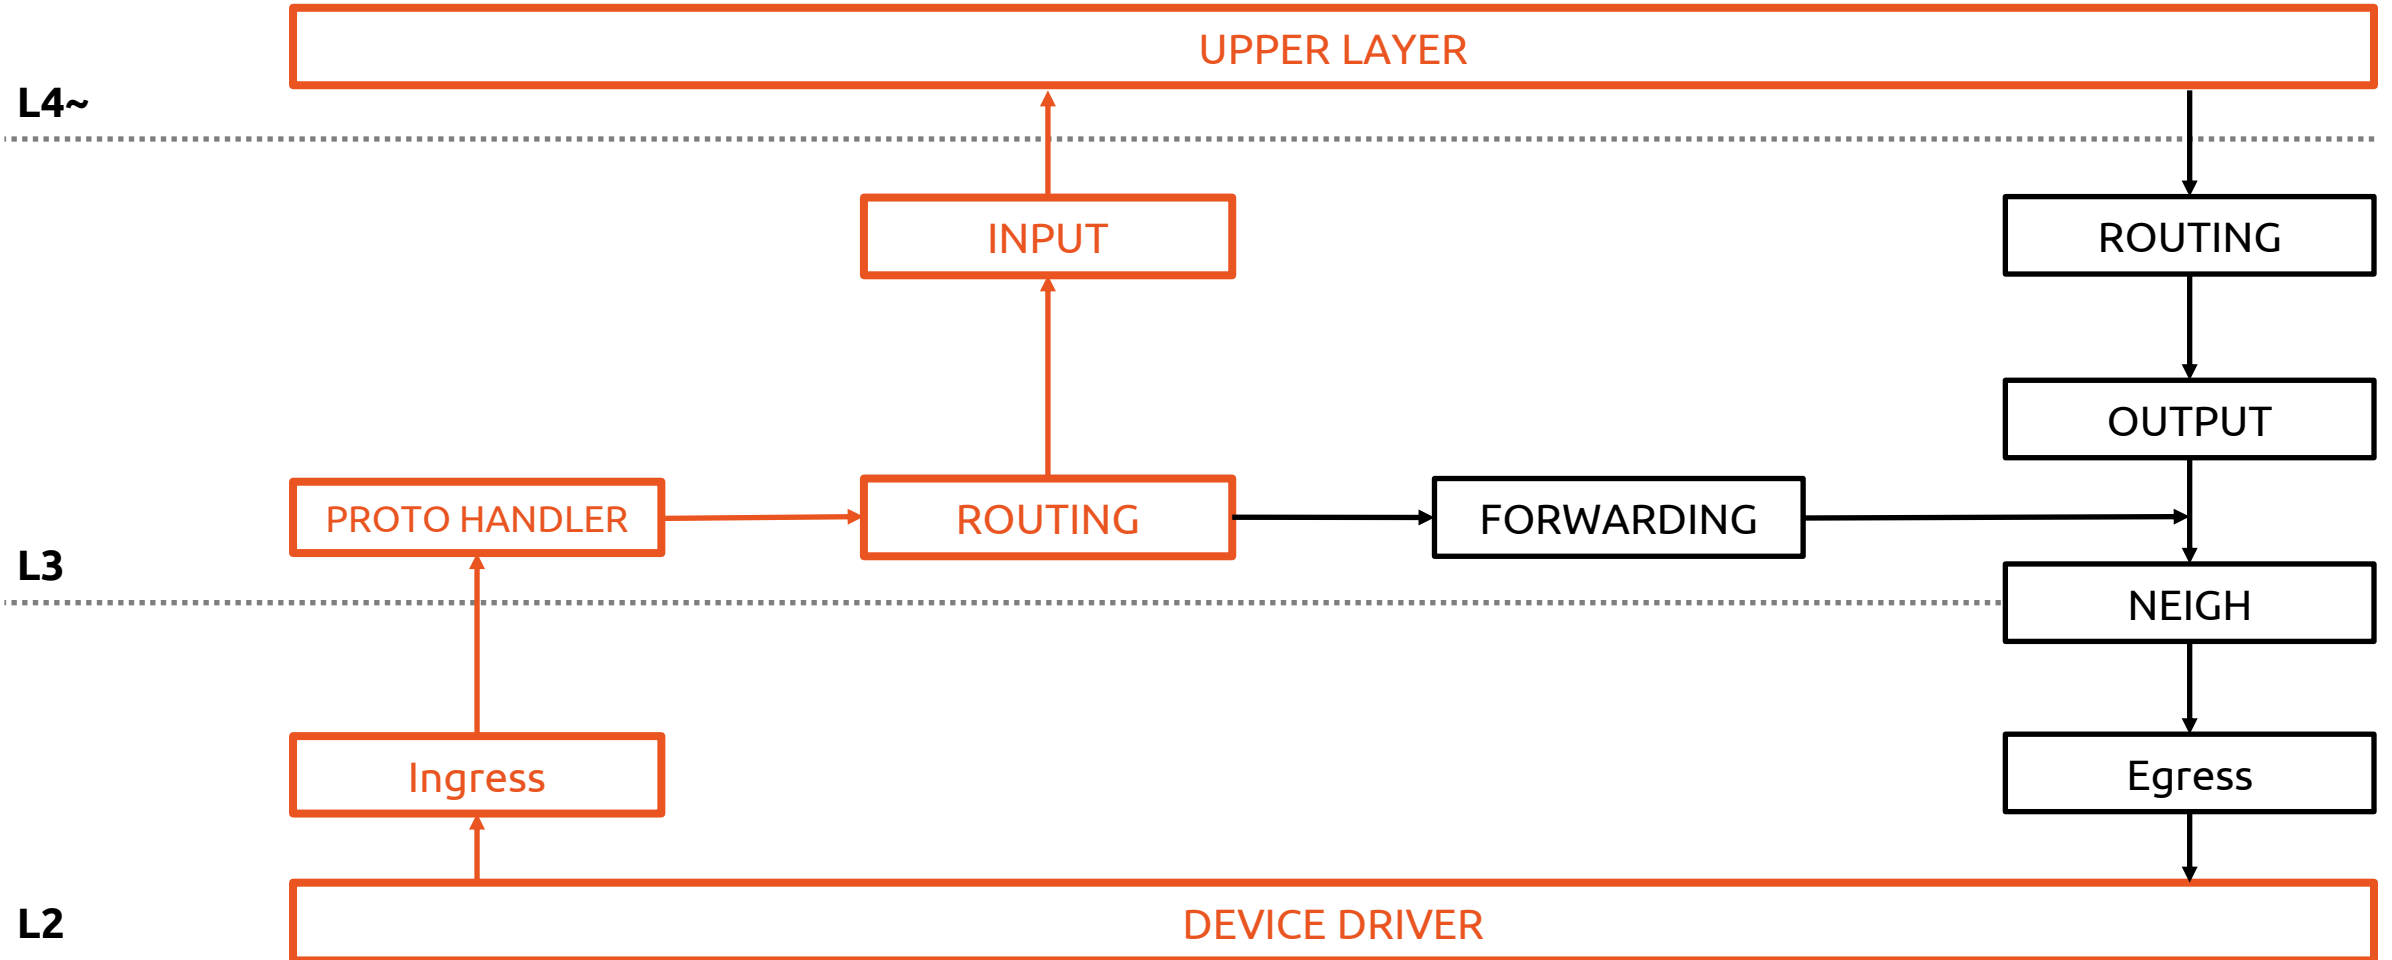


Packet Receiving Path



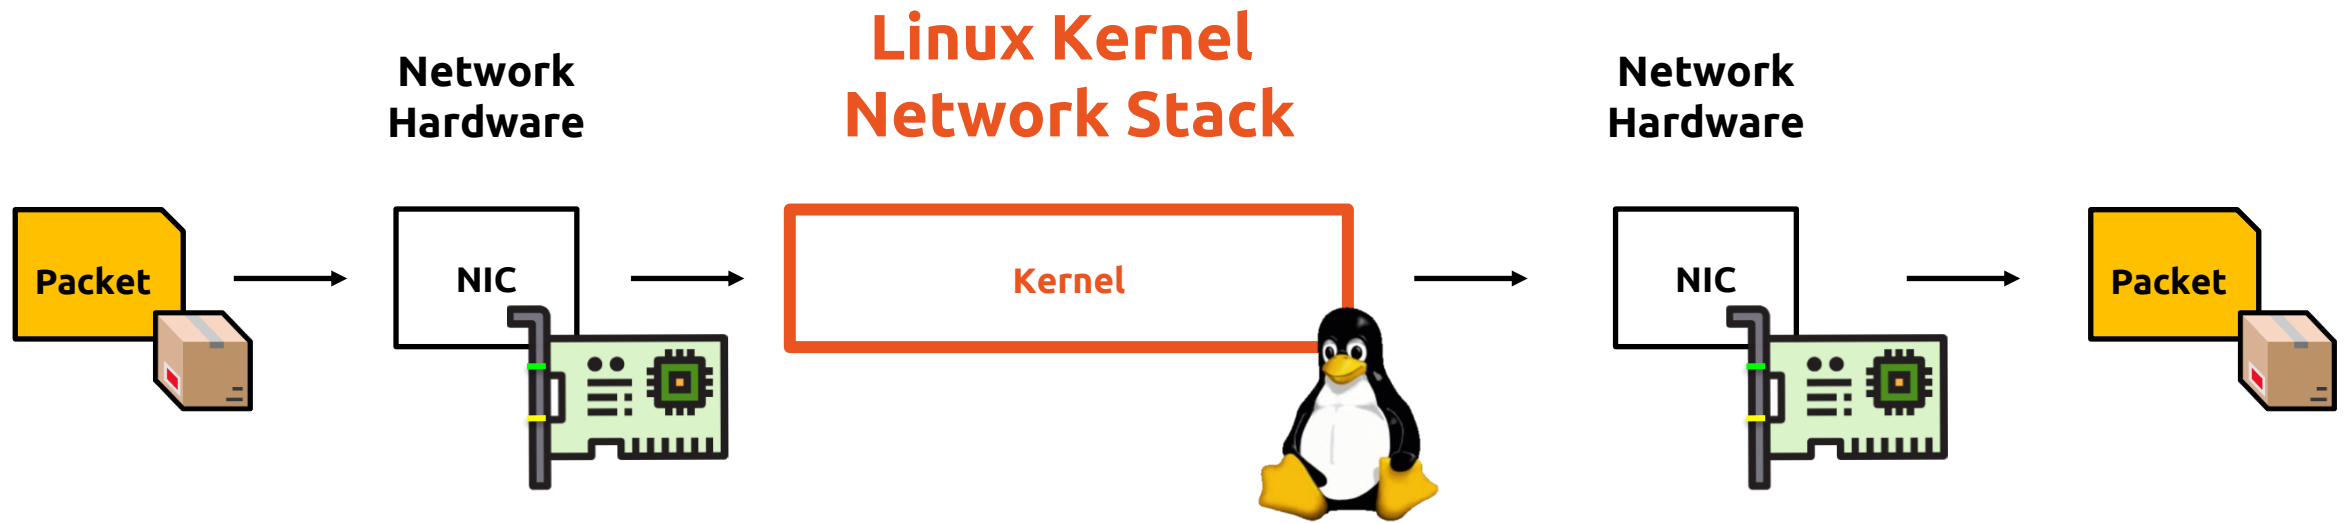


Packet Receiving Path





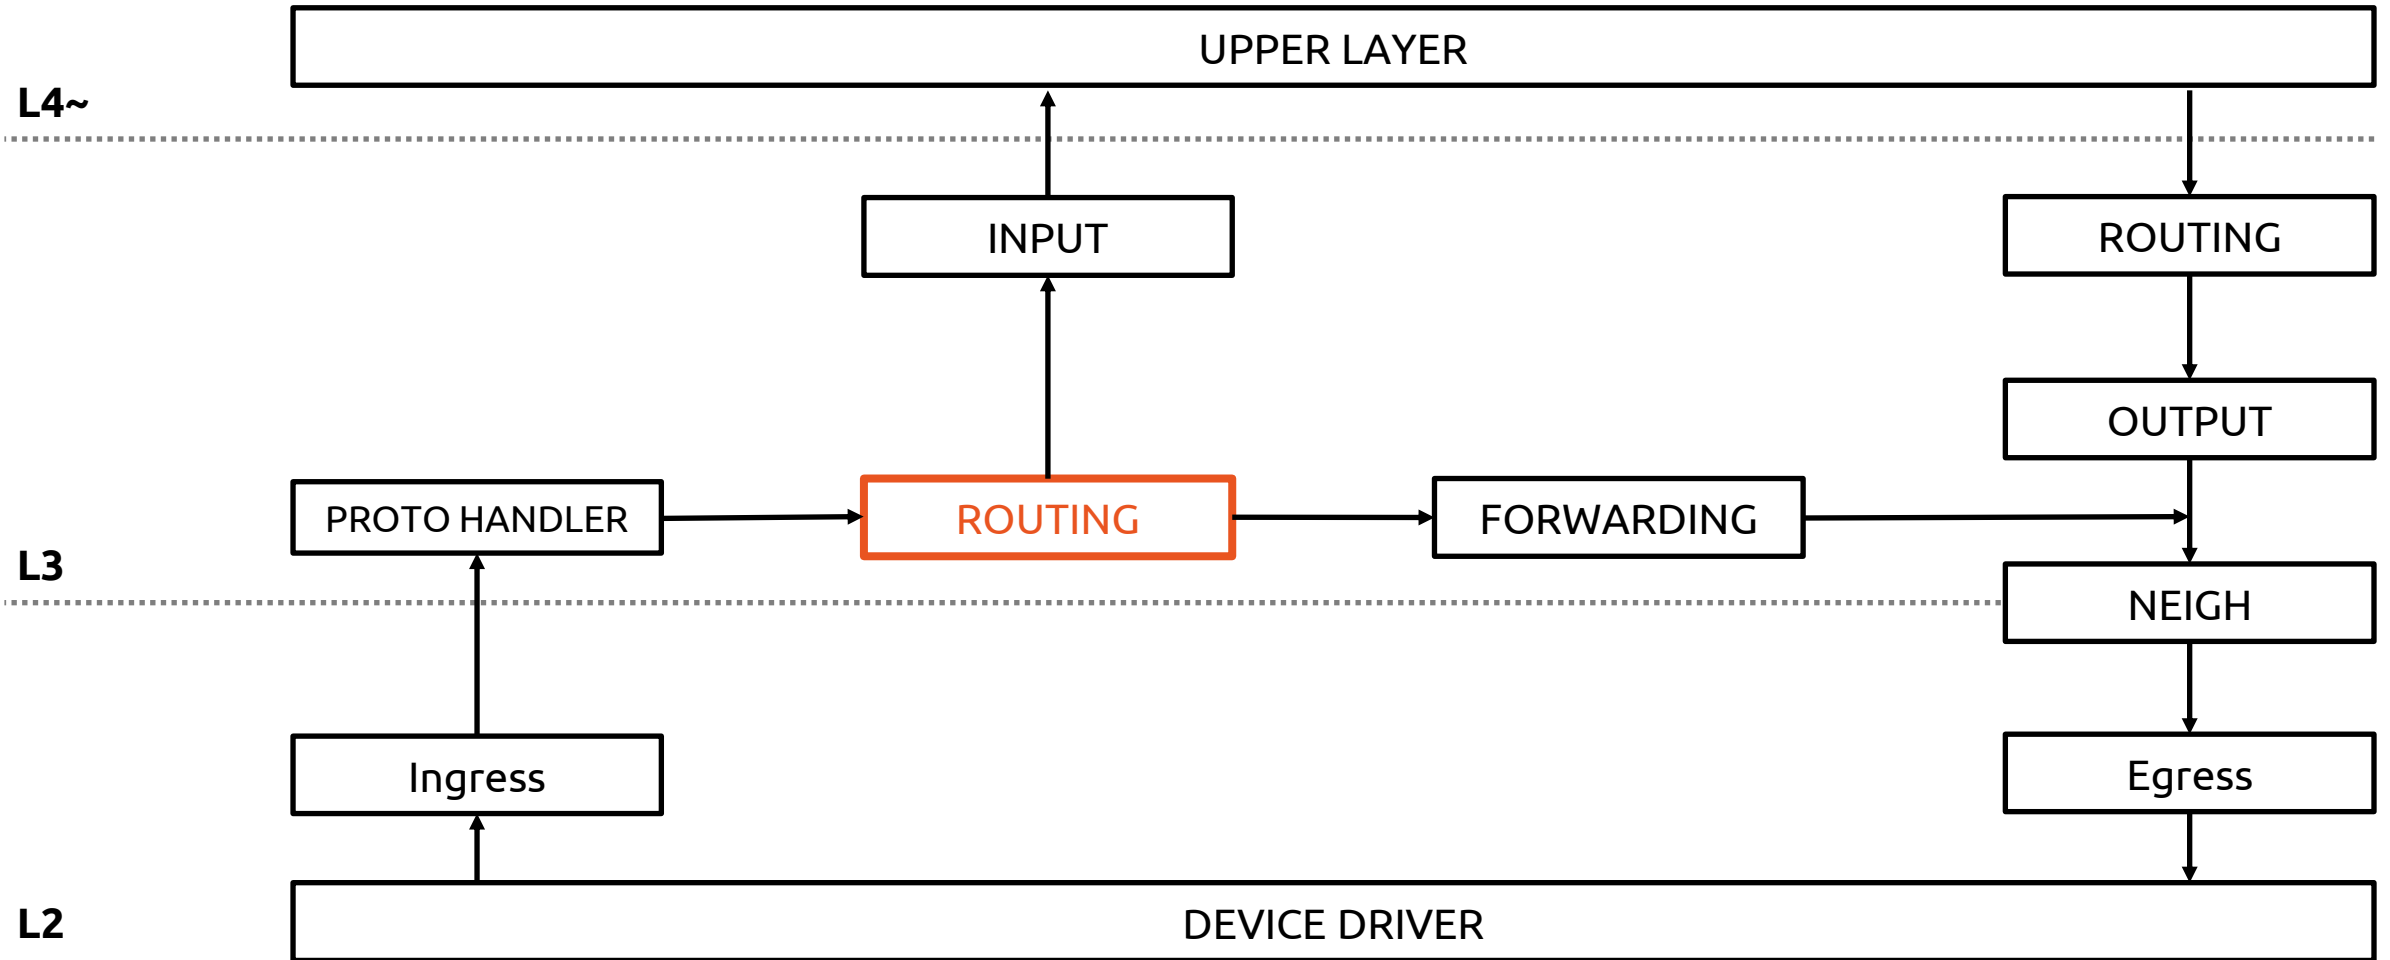
Packet Forwarding Path





Packet Forwarding Path

Lookup the ROUTING table to identify where the packet is headed





Packet Forwarding Path

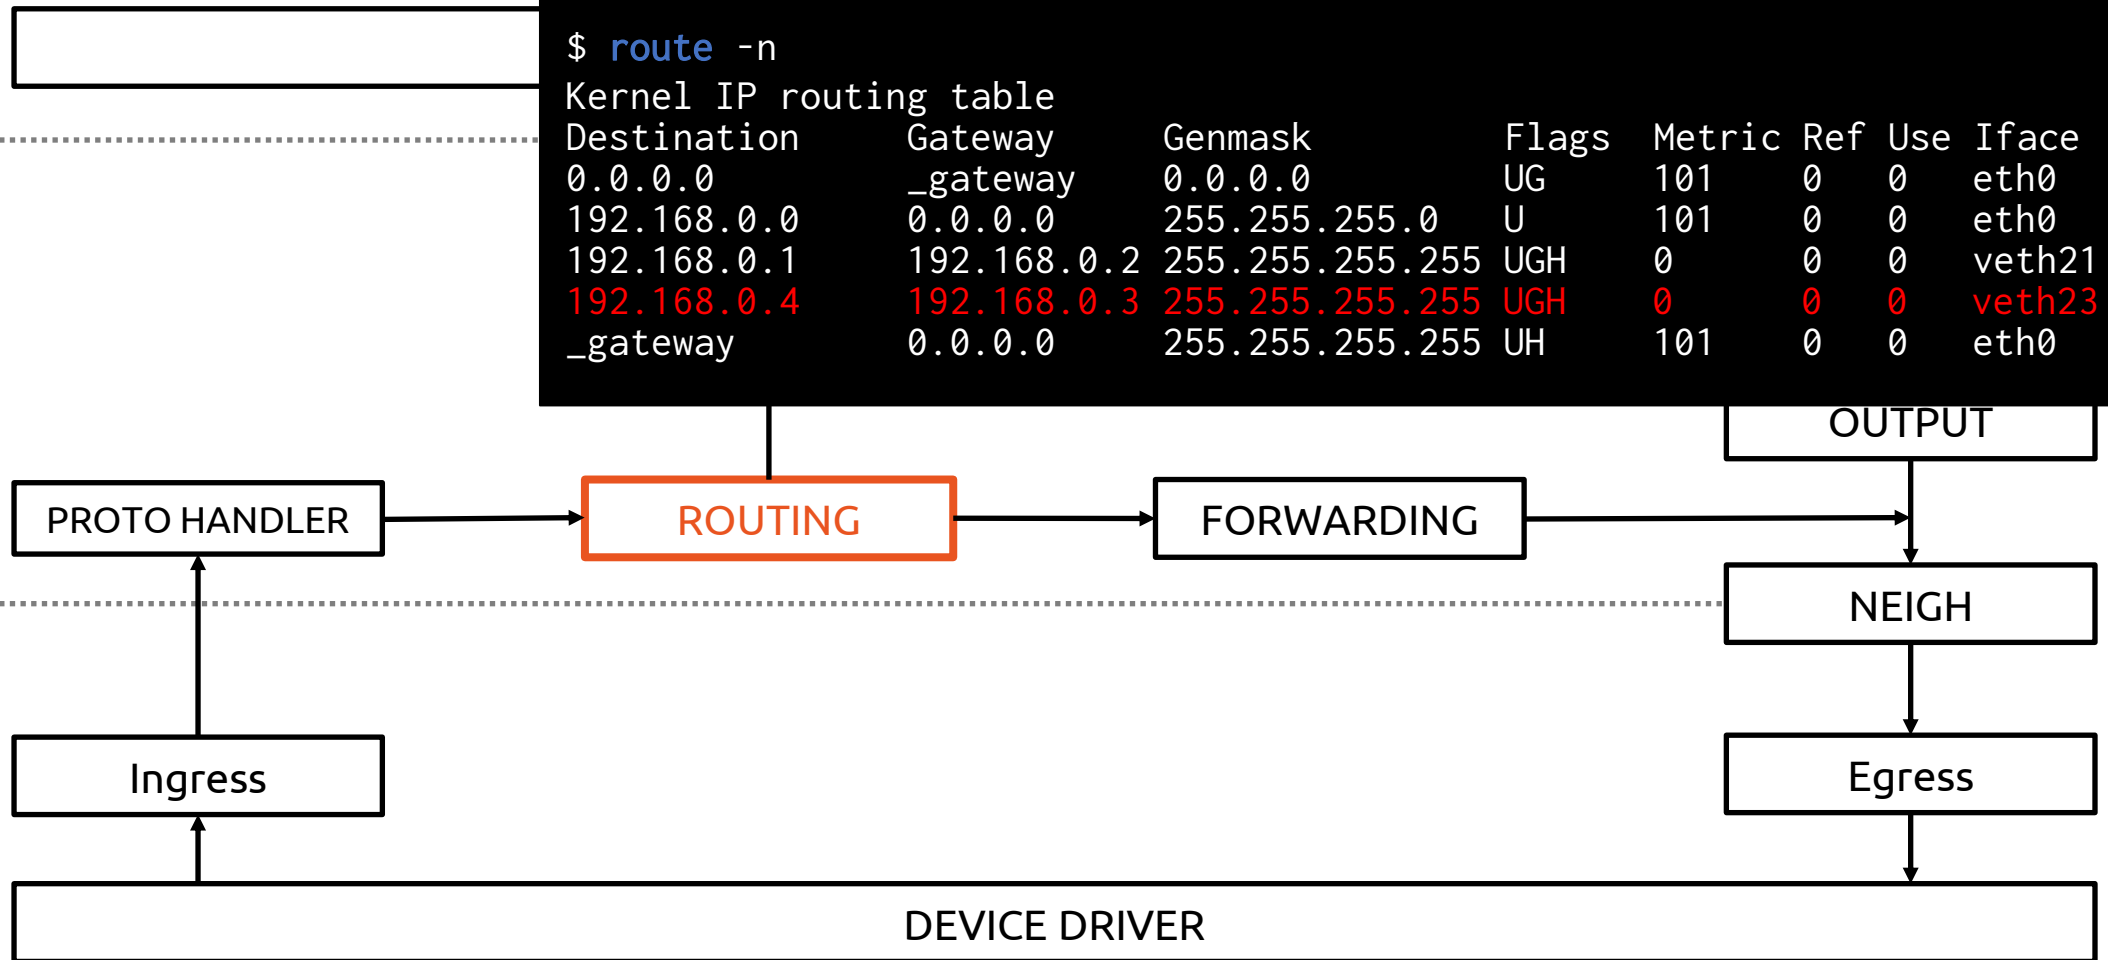
Lookup the ROUTING table to identify where the packet is headed

```
$ route -n
Kernel IP routing table
Destination      Gateway         Genmask         Flags   Metric  Ref  Use  Iface
0.0.0.0          _gateway       0.0.0.0         UG      101     0   0   eth0
192.168.0.0      0.0.0.0        255.255.255.0   U       101     0   0   eth0
192.168.0.1      192.168.0.2    255.255.255.255 UGH     0       0   0   veth21
192.168.0.4      192.168.0.3    255.255.255.255 UGH     0       0   0   veth23
_gateway         0.0.0.0        255.255.255.255 UH      101     0   0   eth0
```

L4~

L3

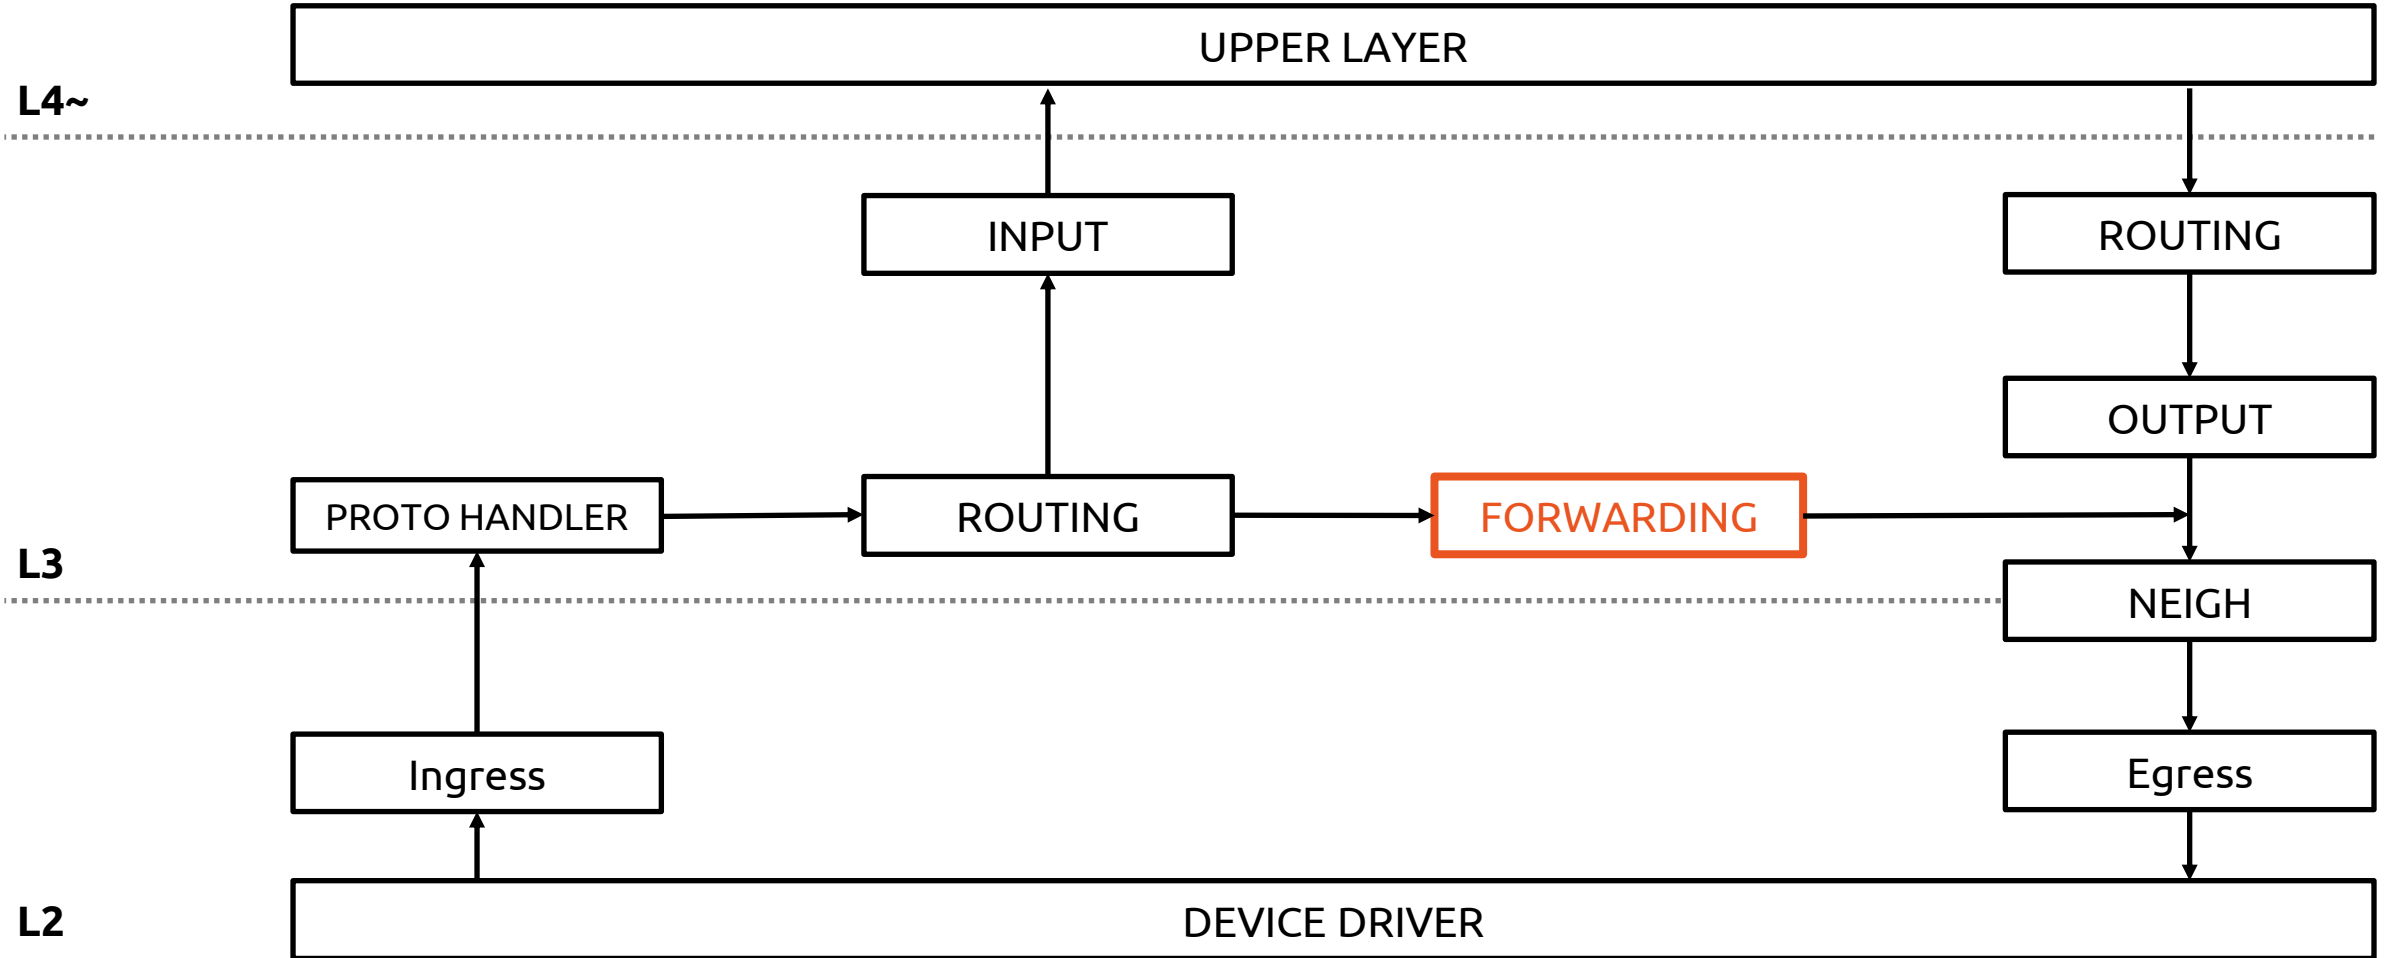
L2





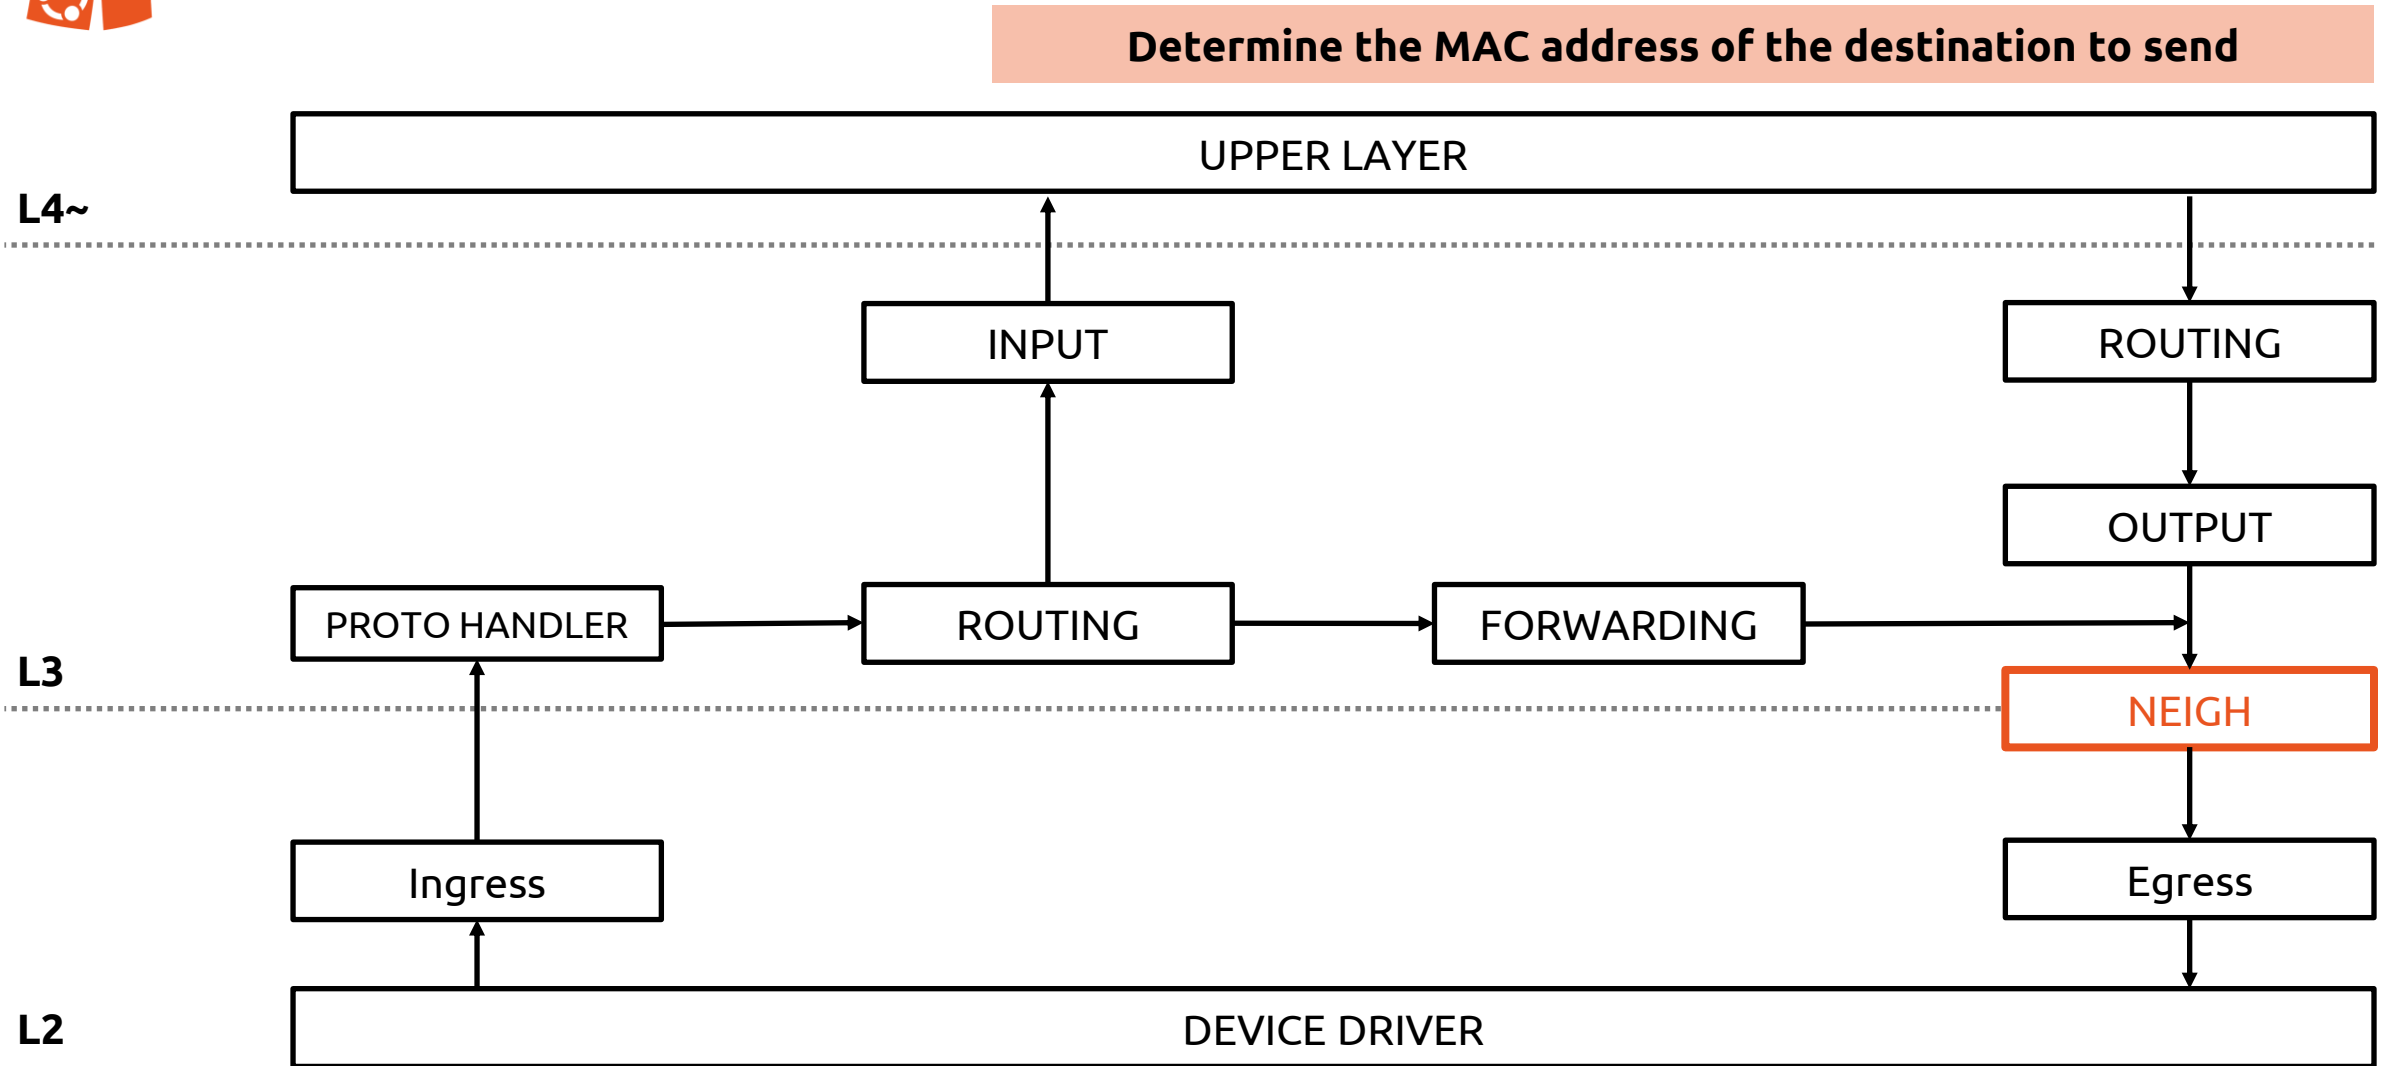
Packet Forwarding Path

If forwarding is required, send to corresponding interface



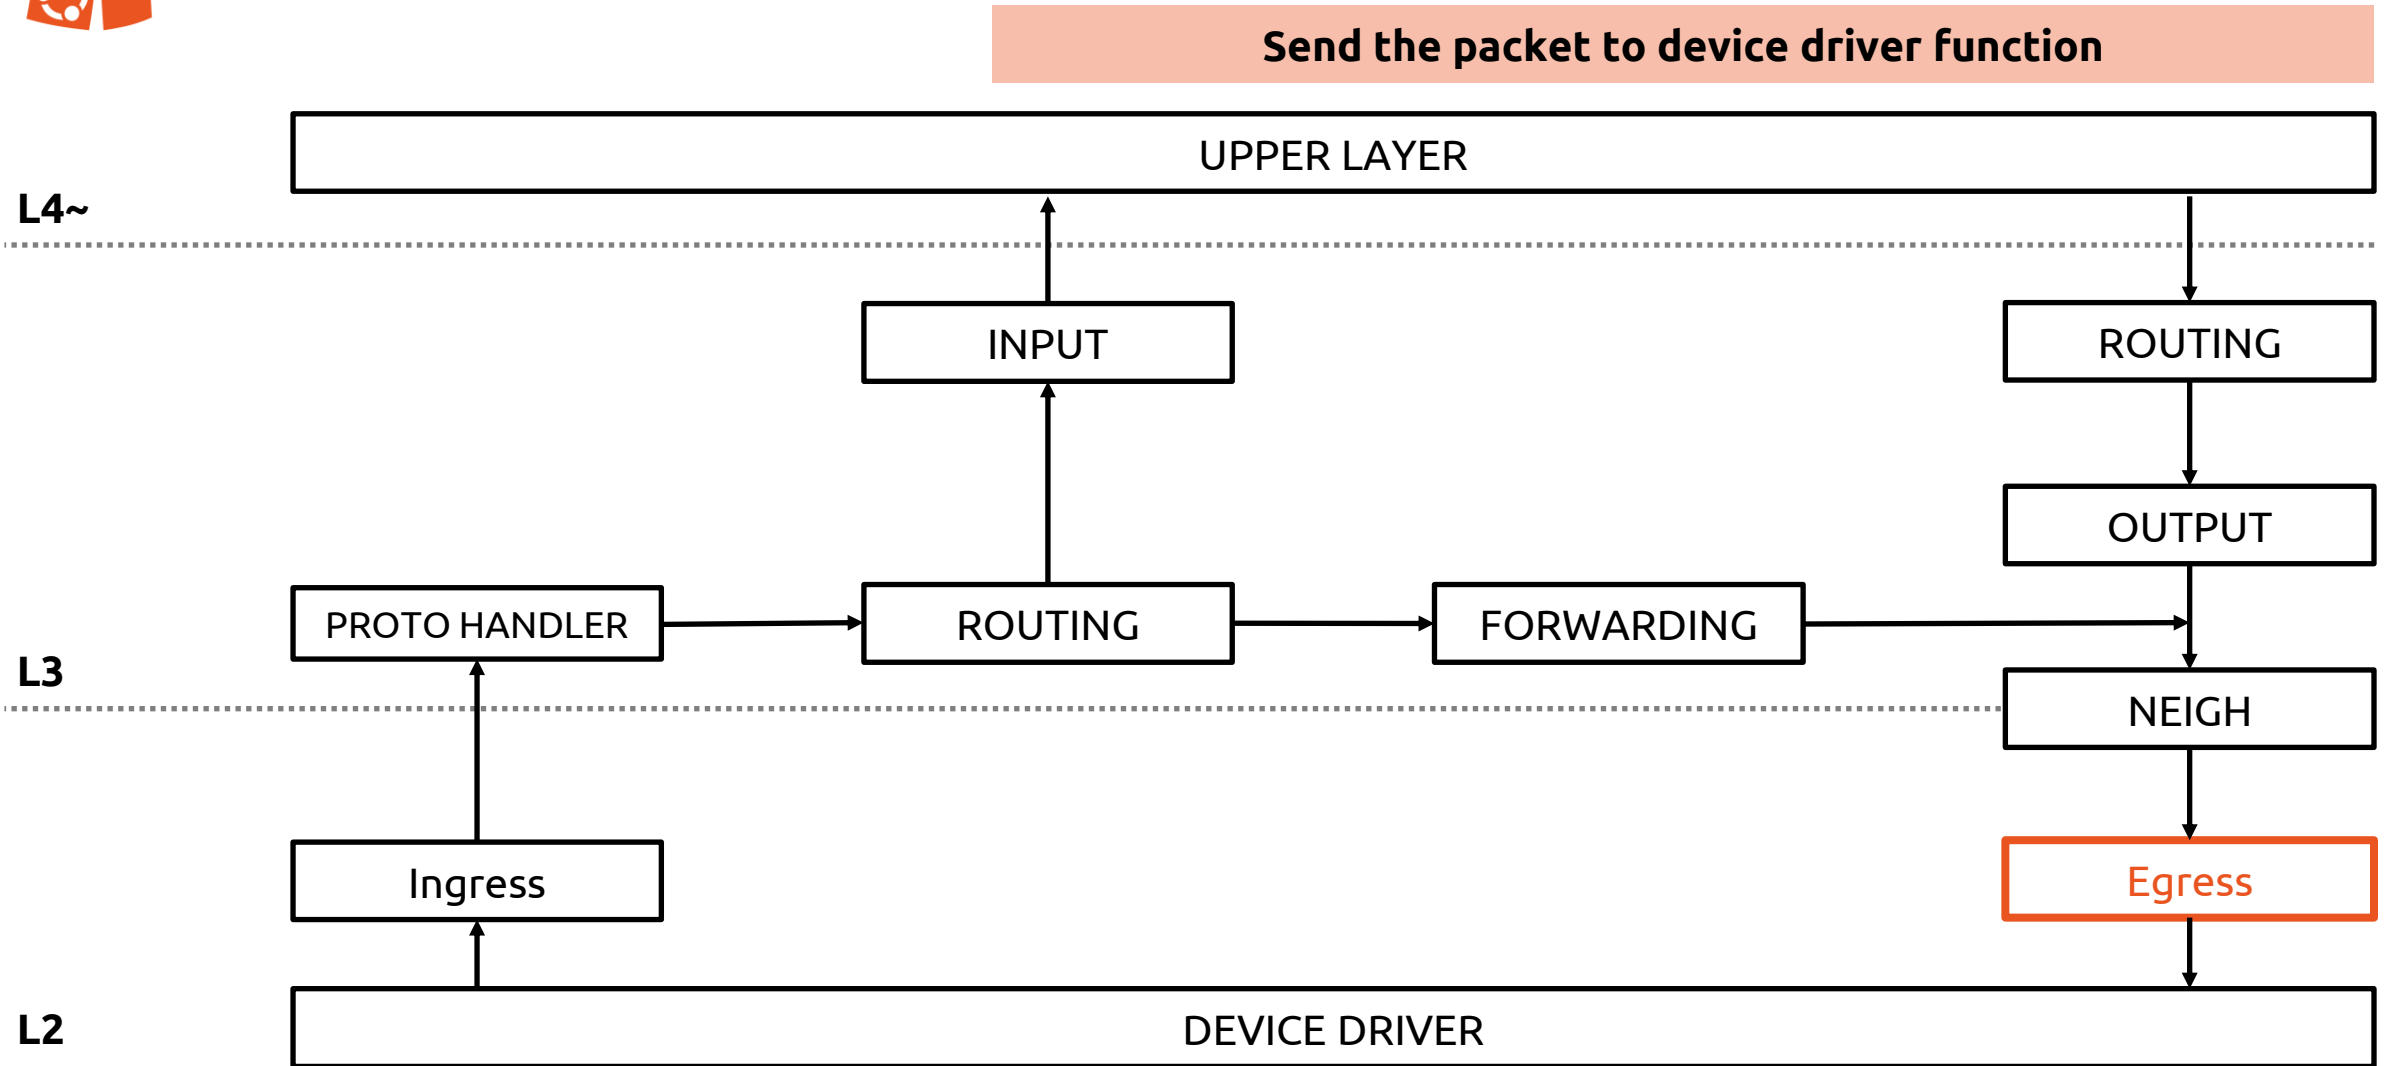


Packet Forwarding Path



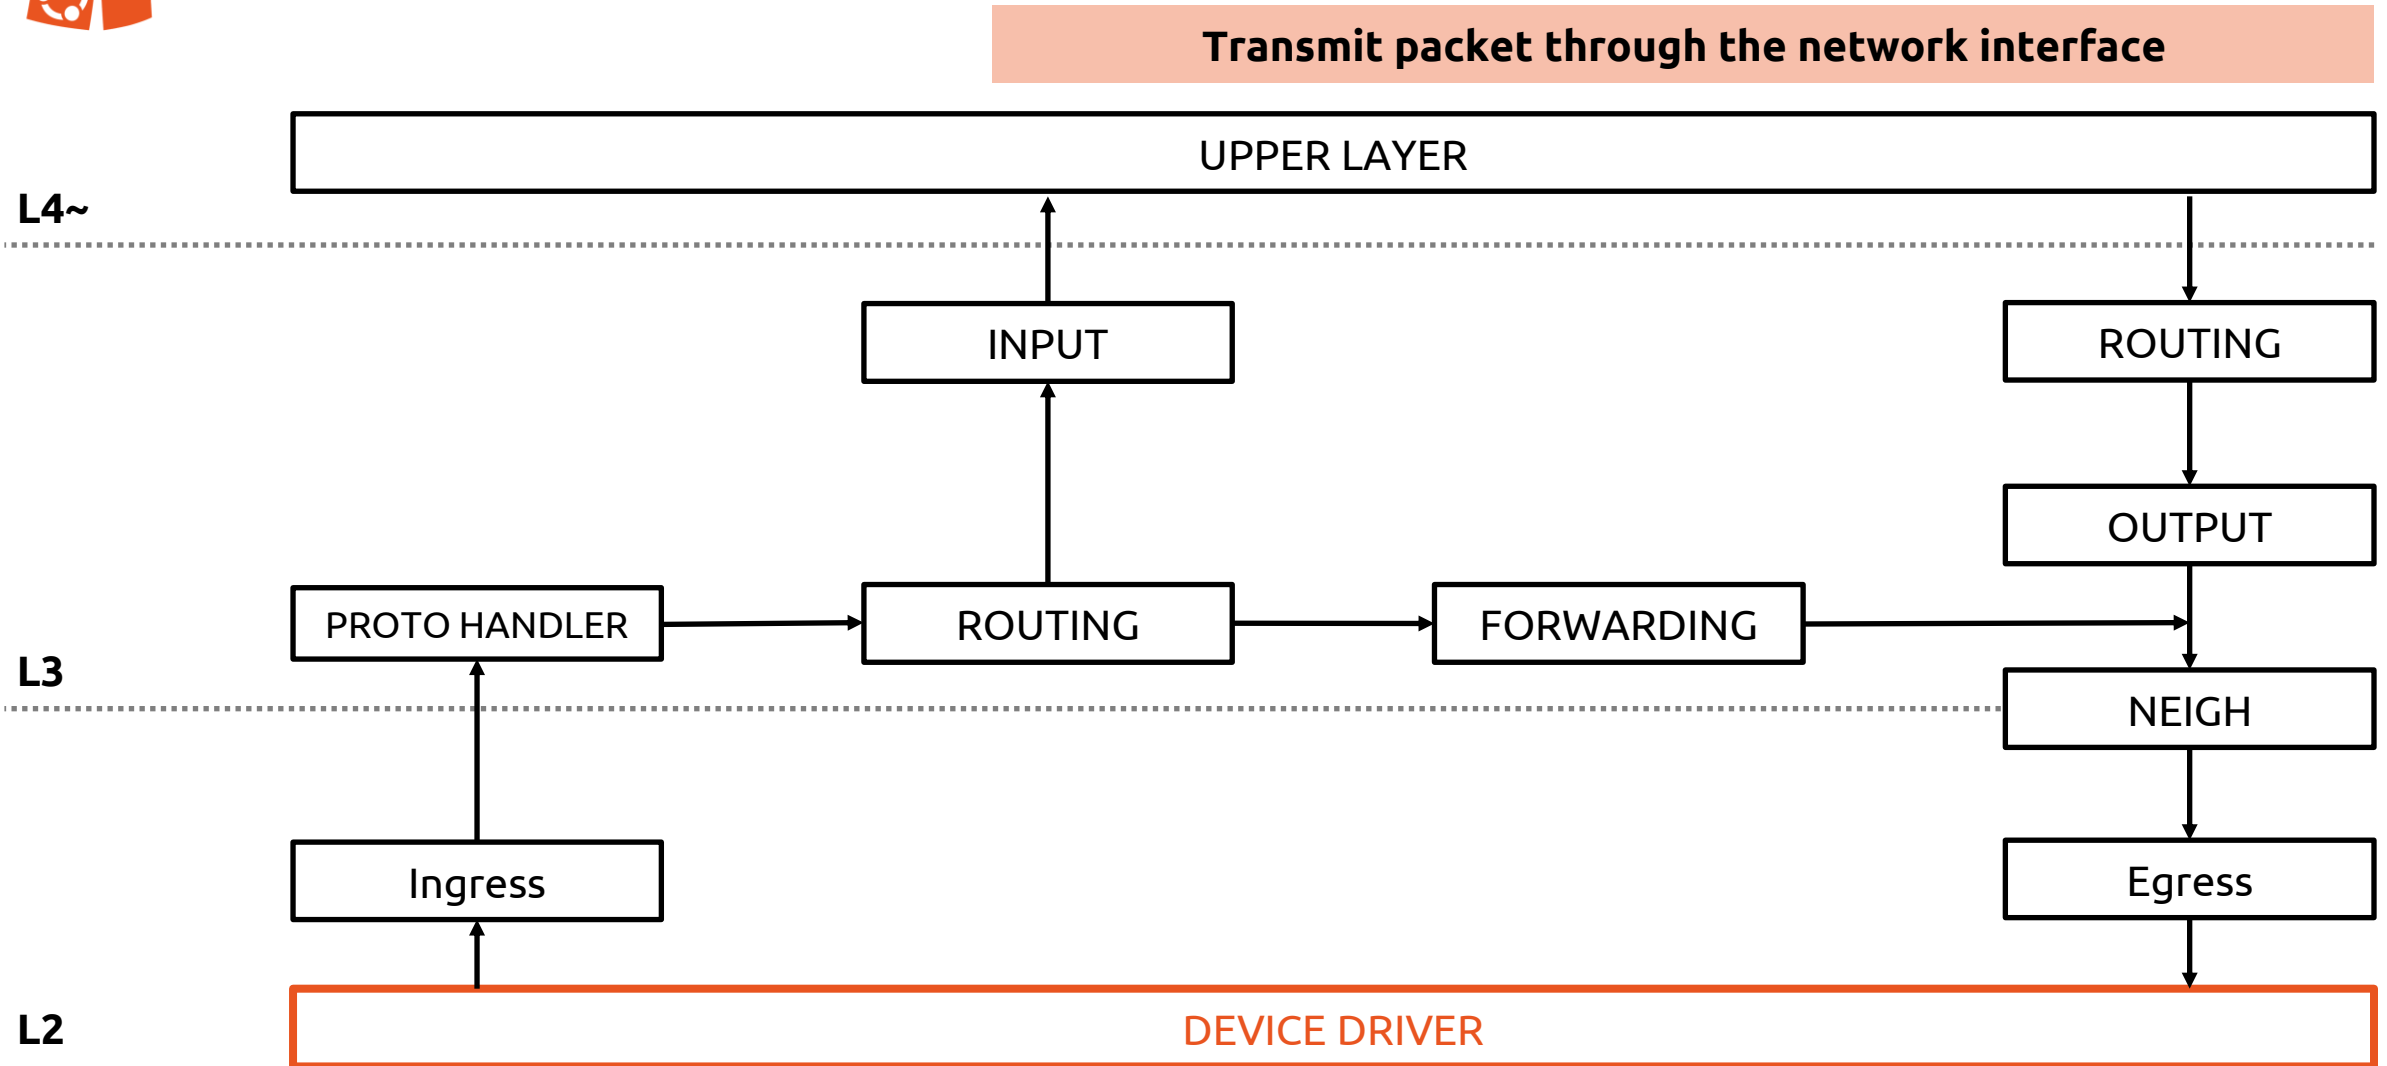


Packet Forwarding Path



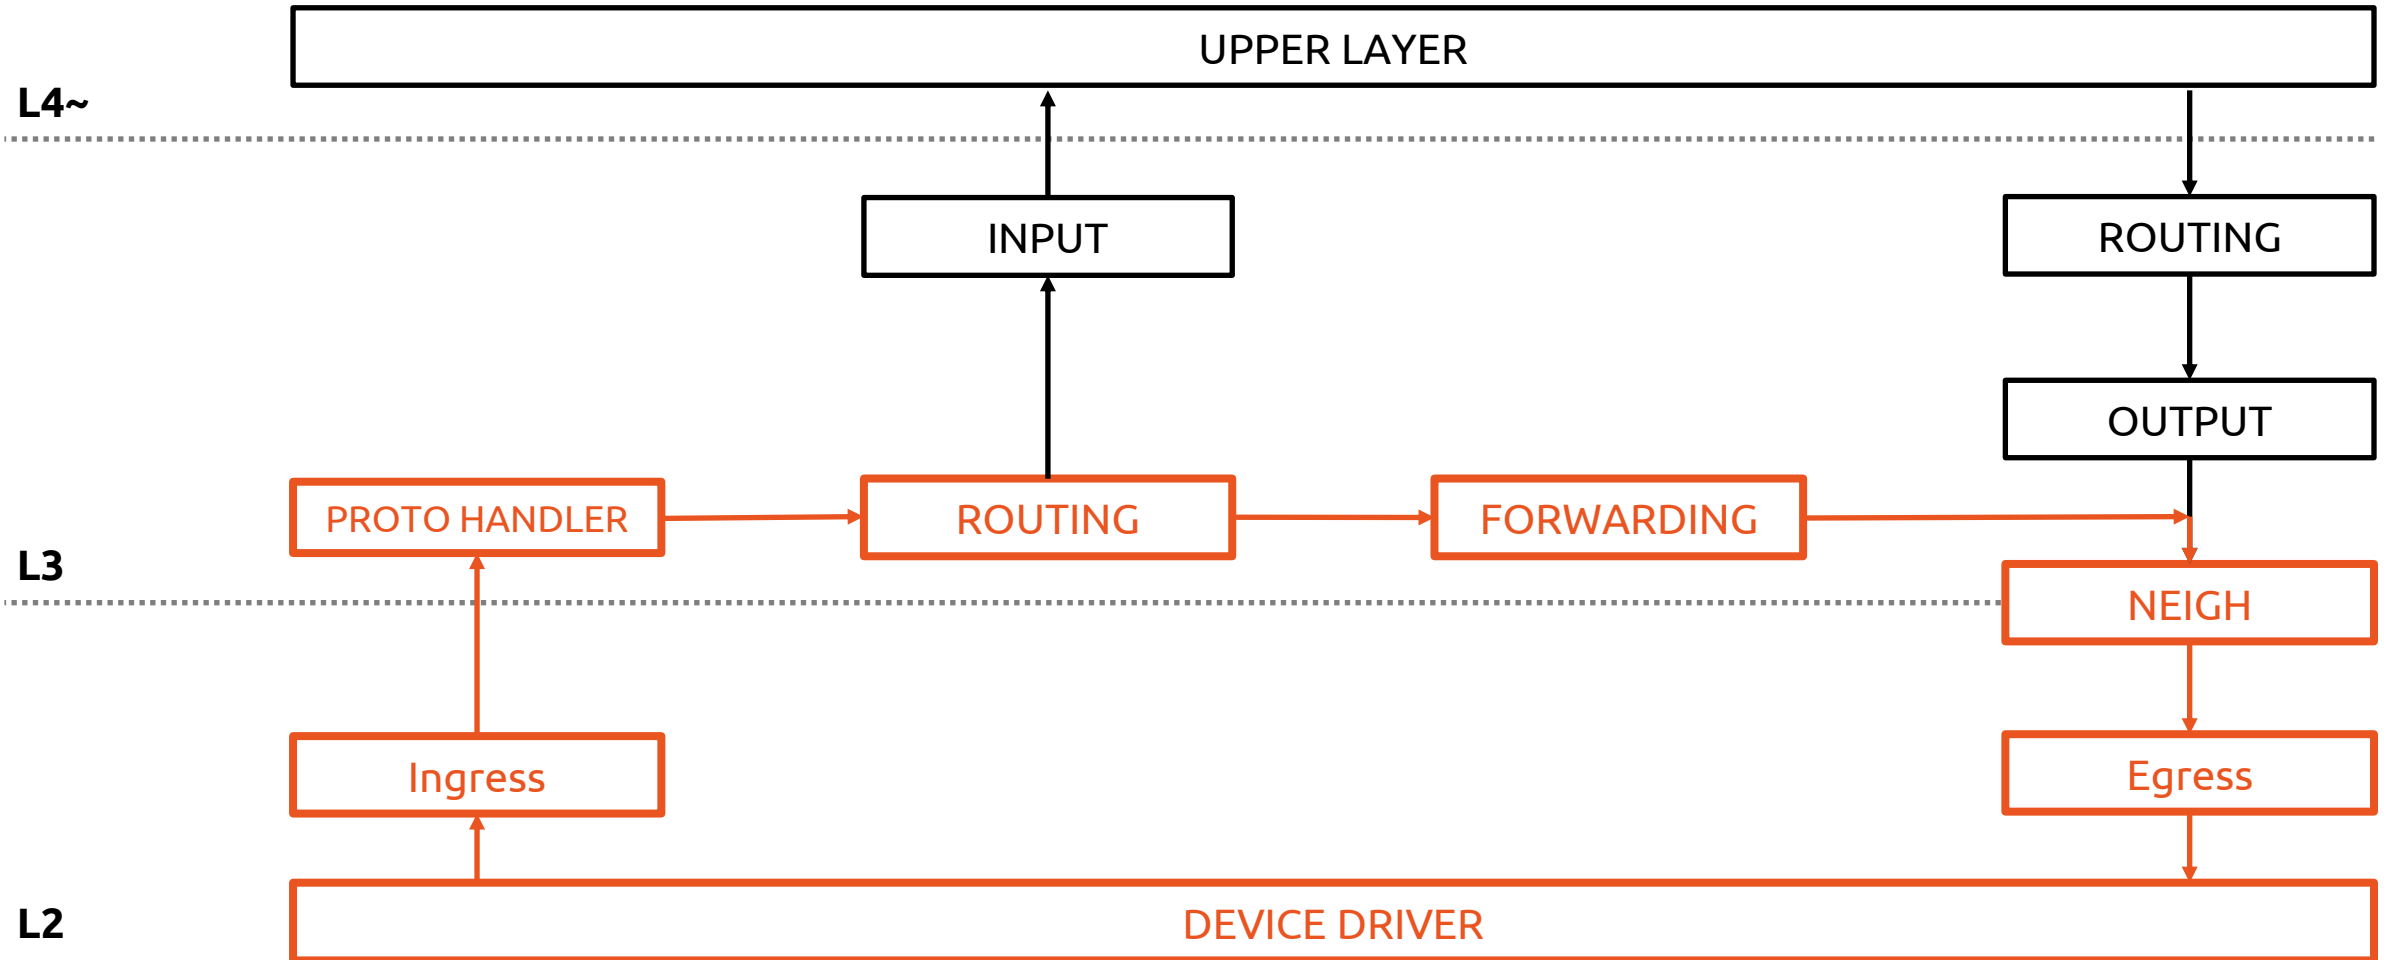


Packet Forwarding Path



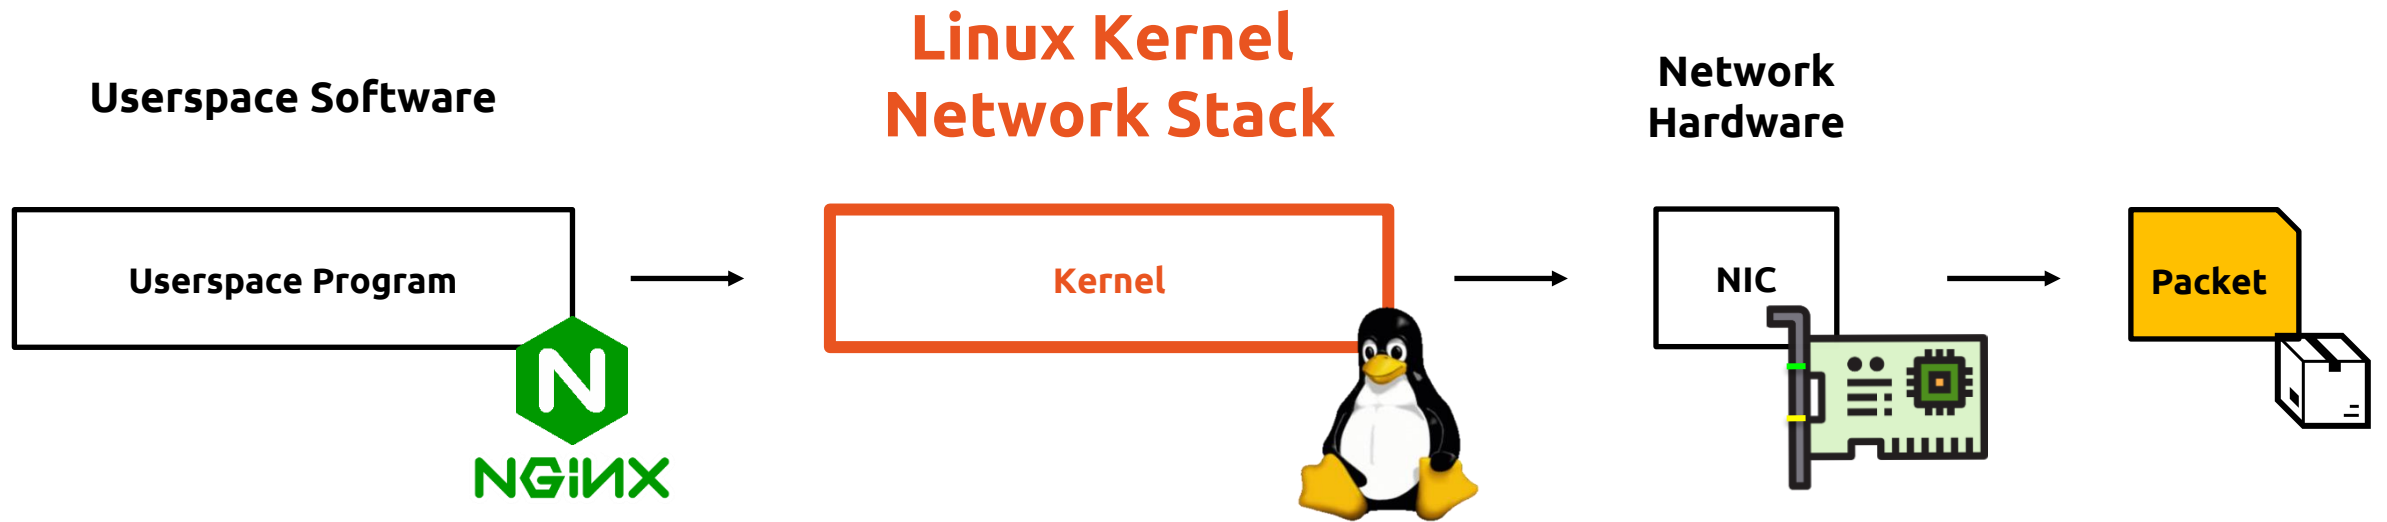


Packet Forwarding Path



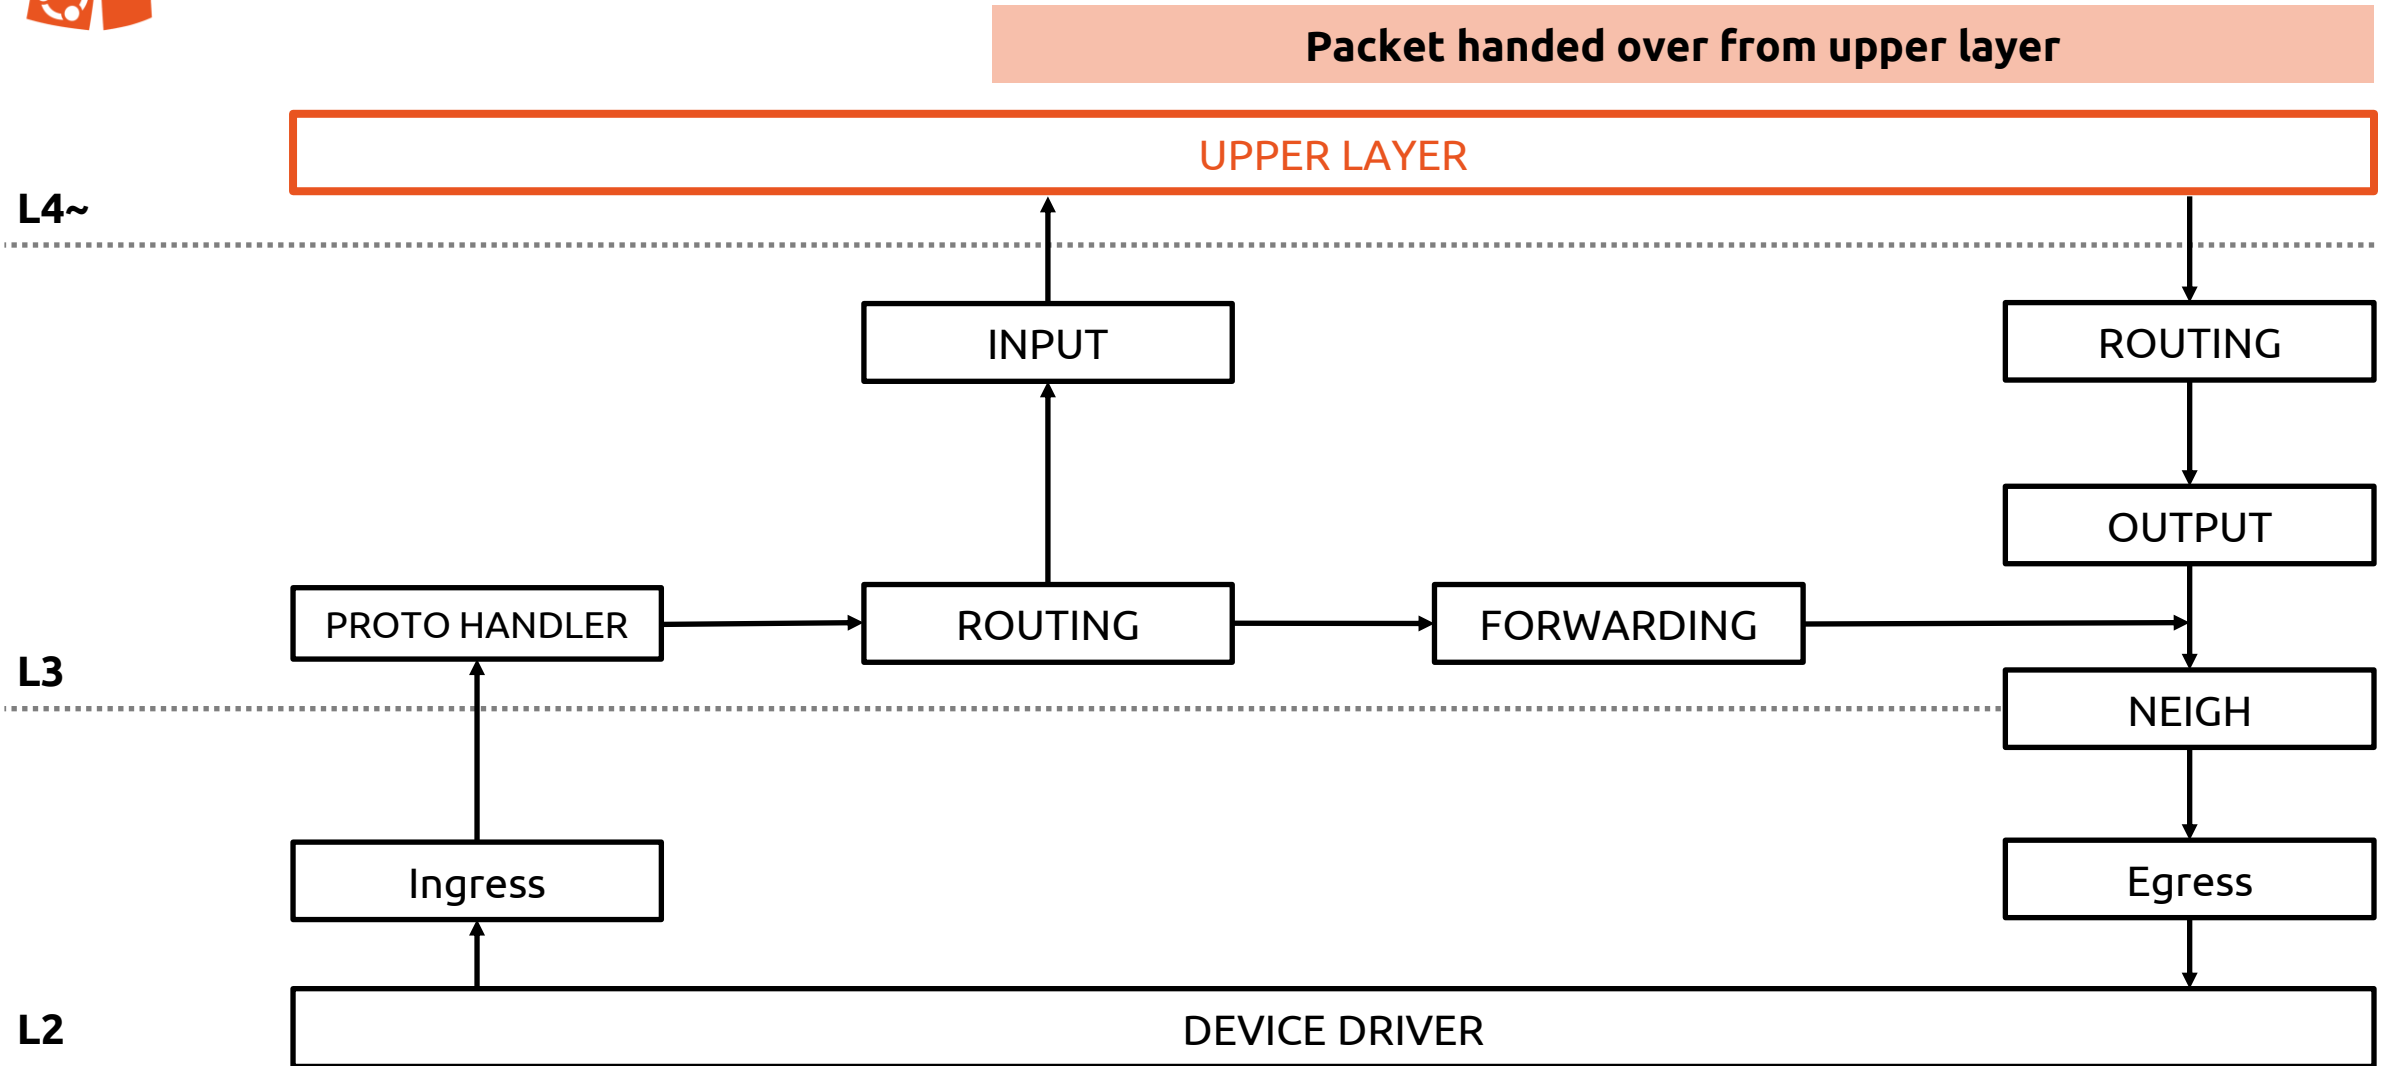


Packet Sending Path






Packet Sending Path





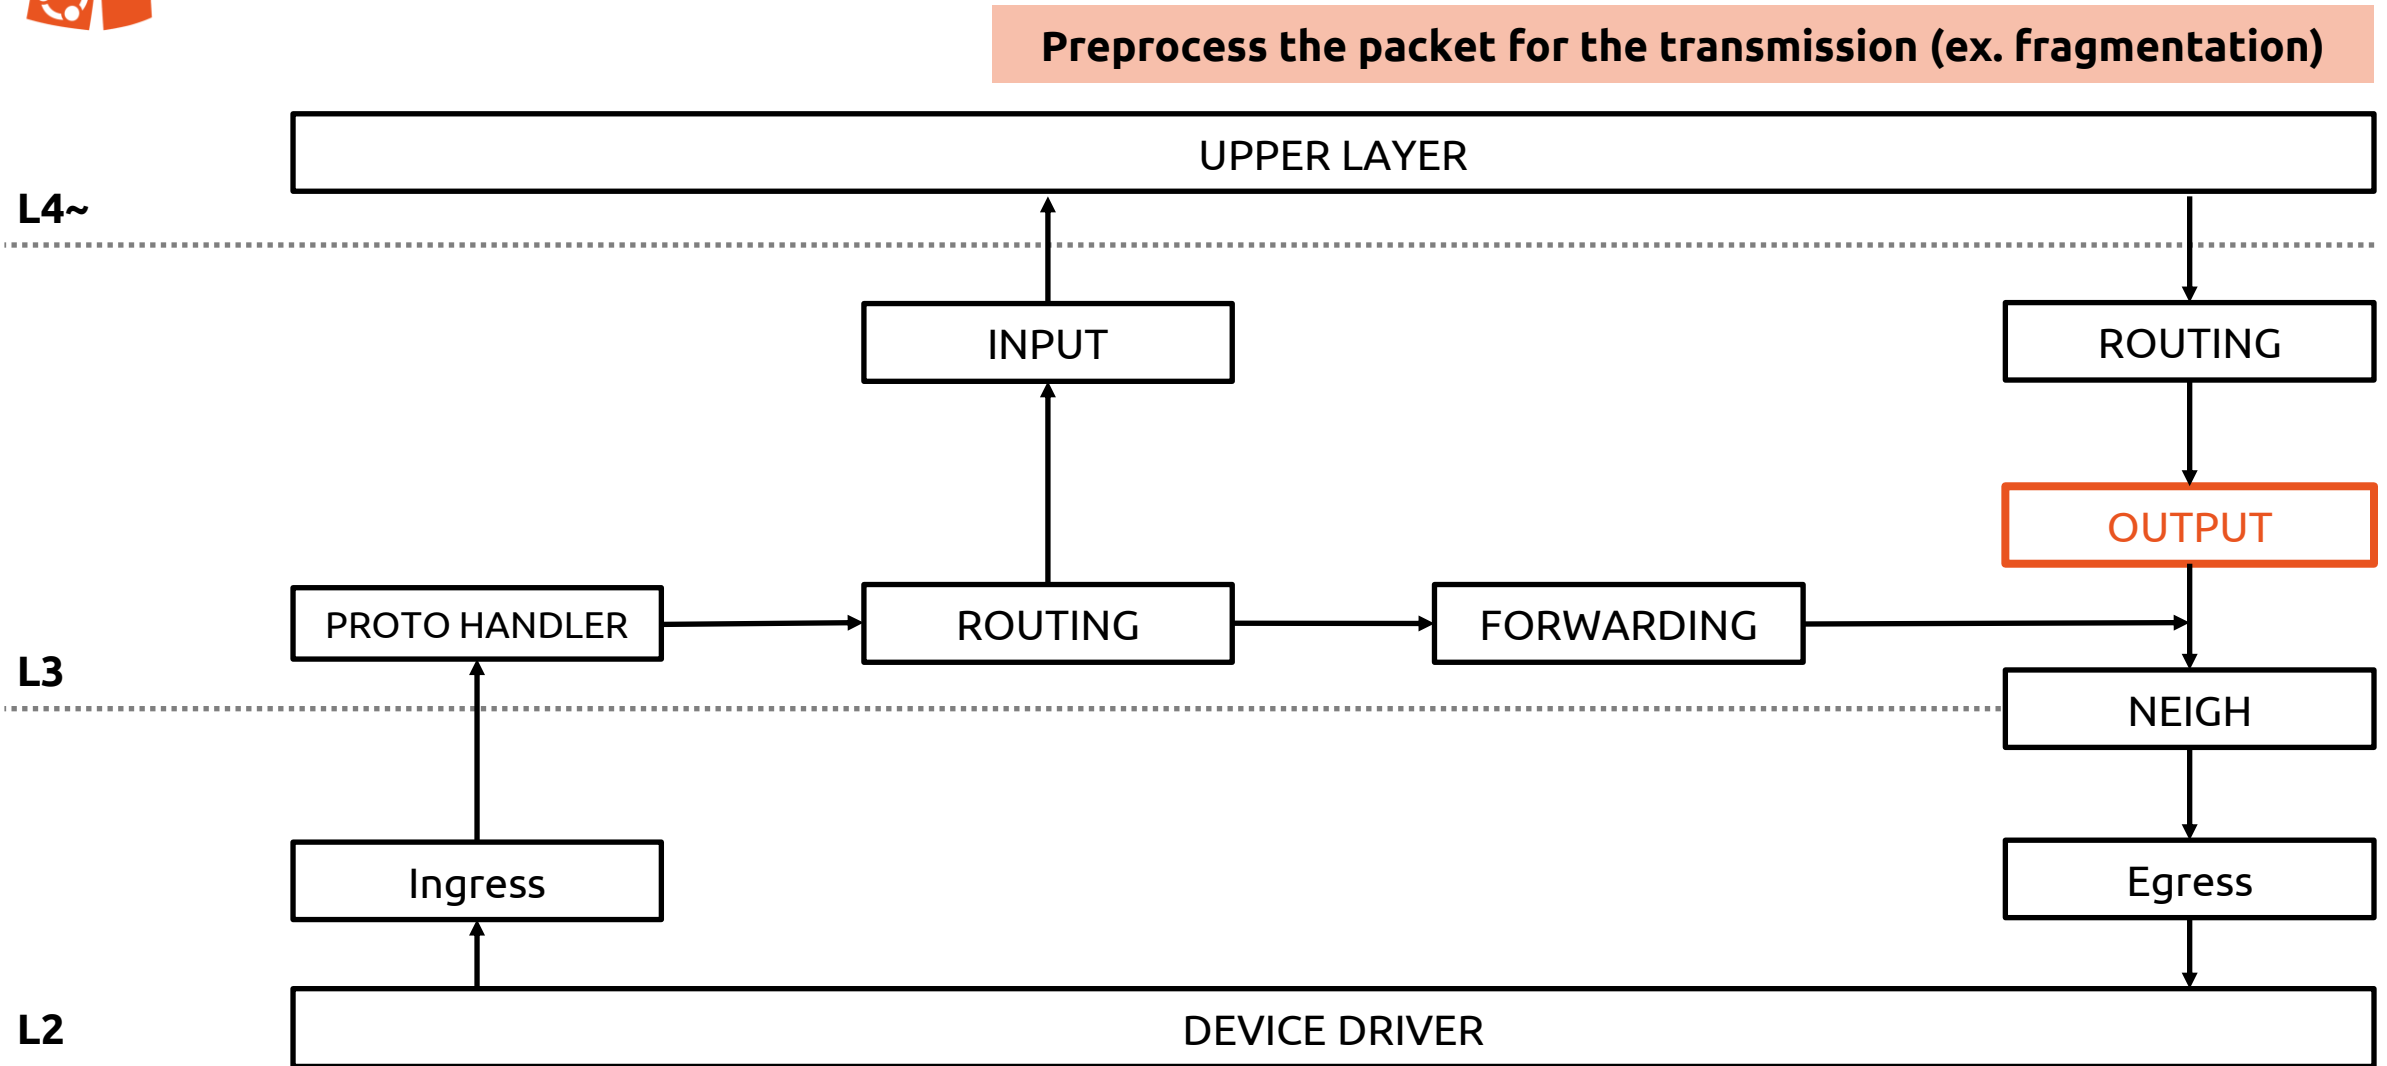
Packet Sending Path

Check the ROUTING table to see which interface to send packet



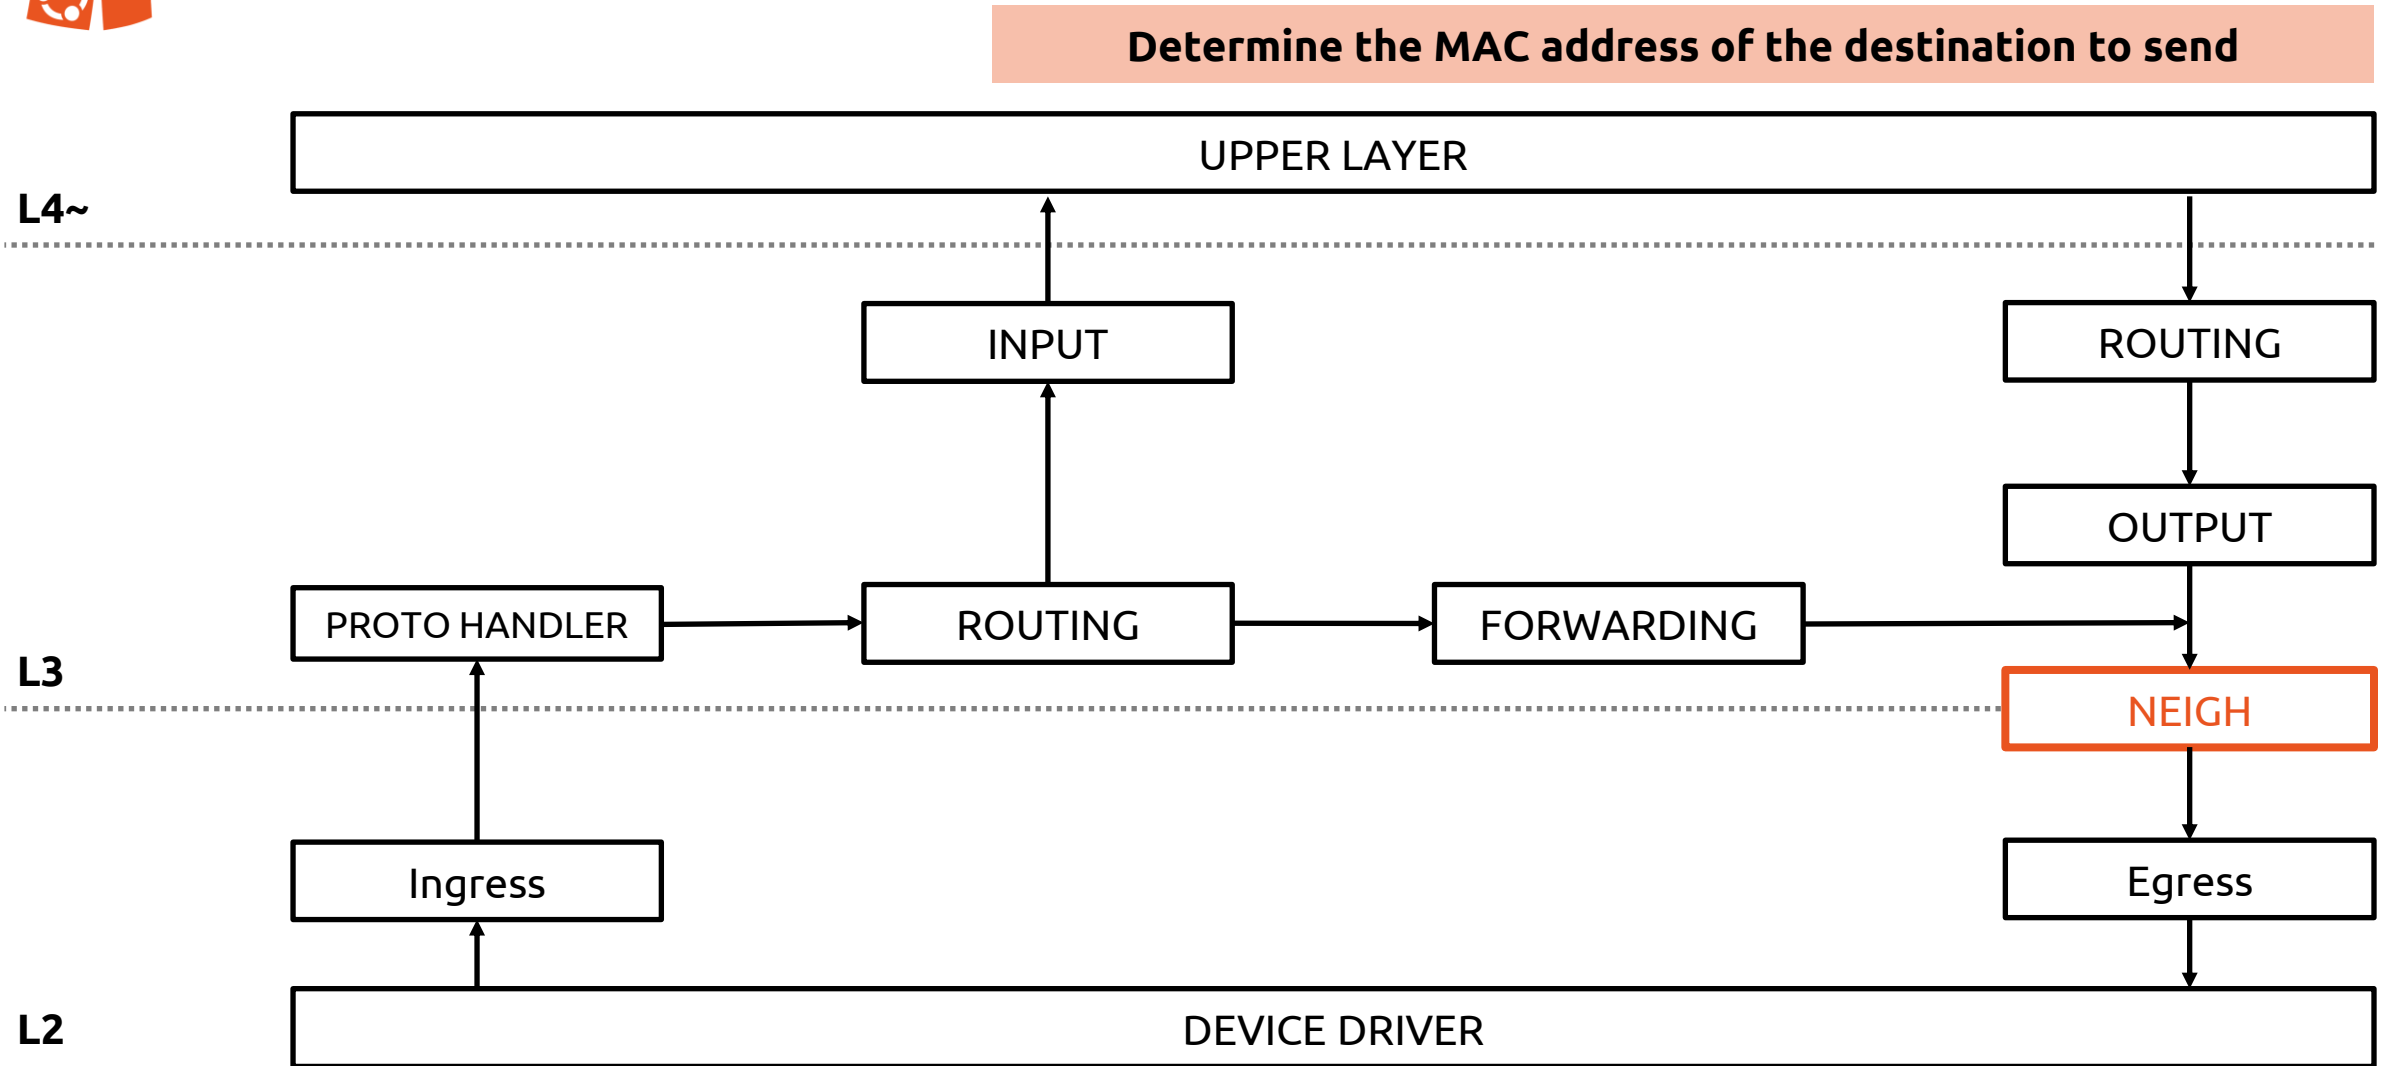


Packet Sending Path



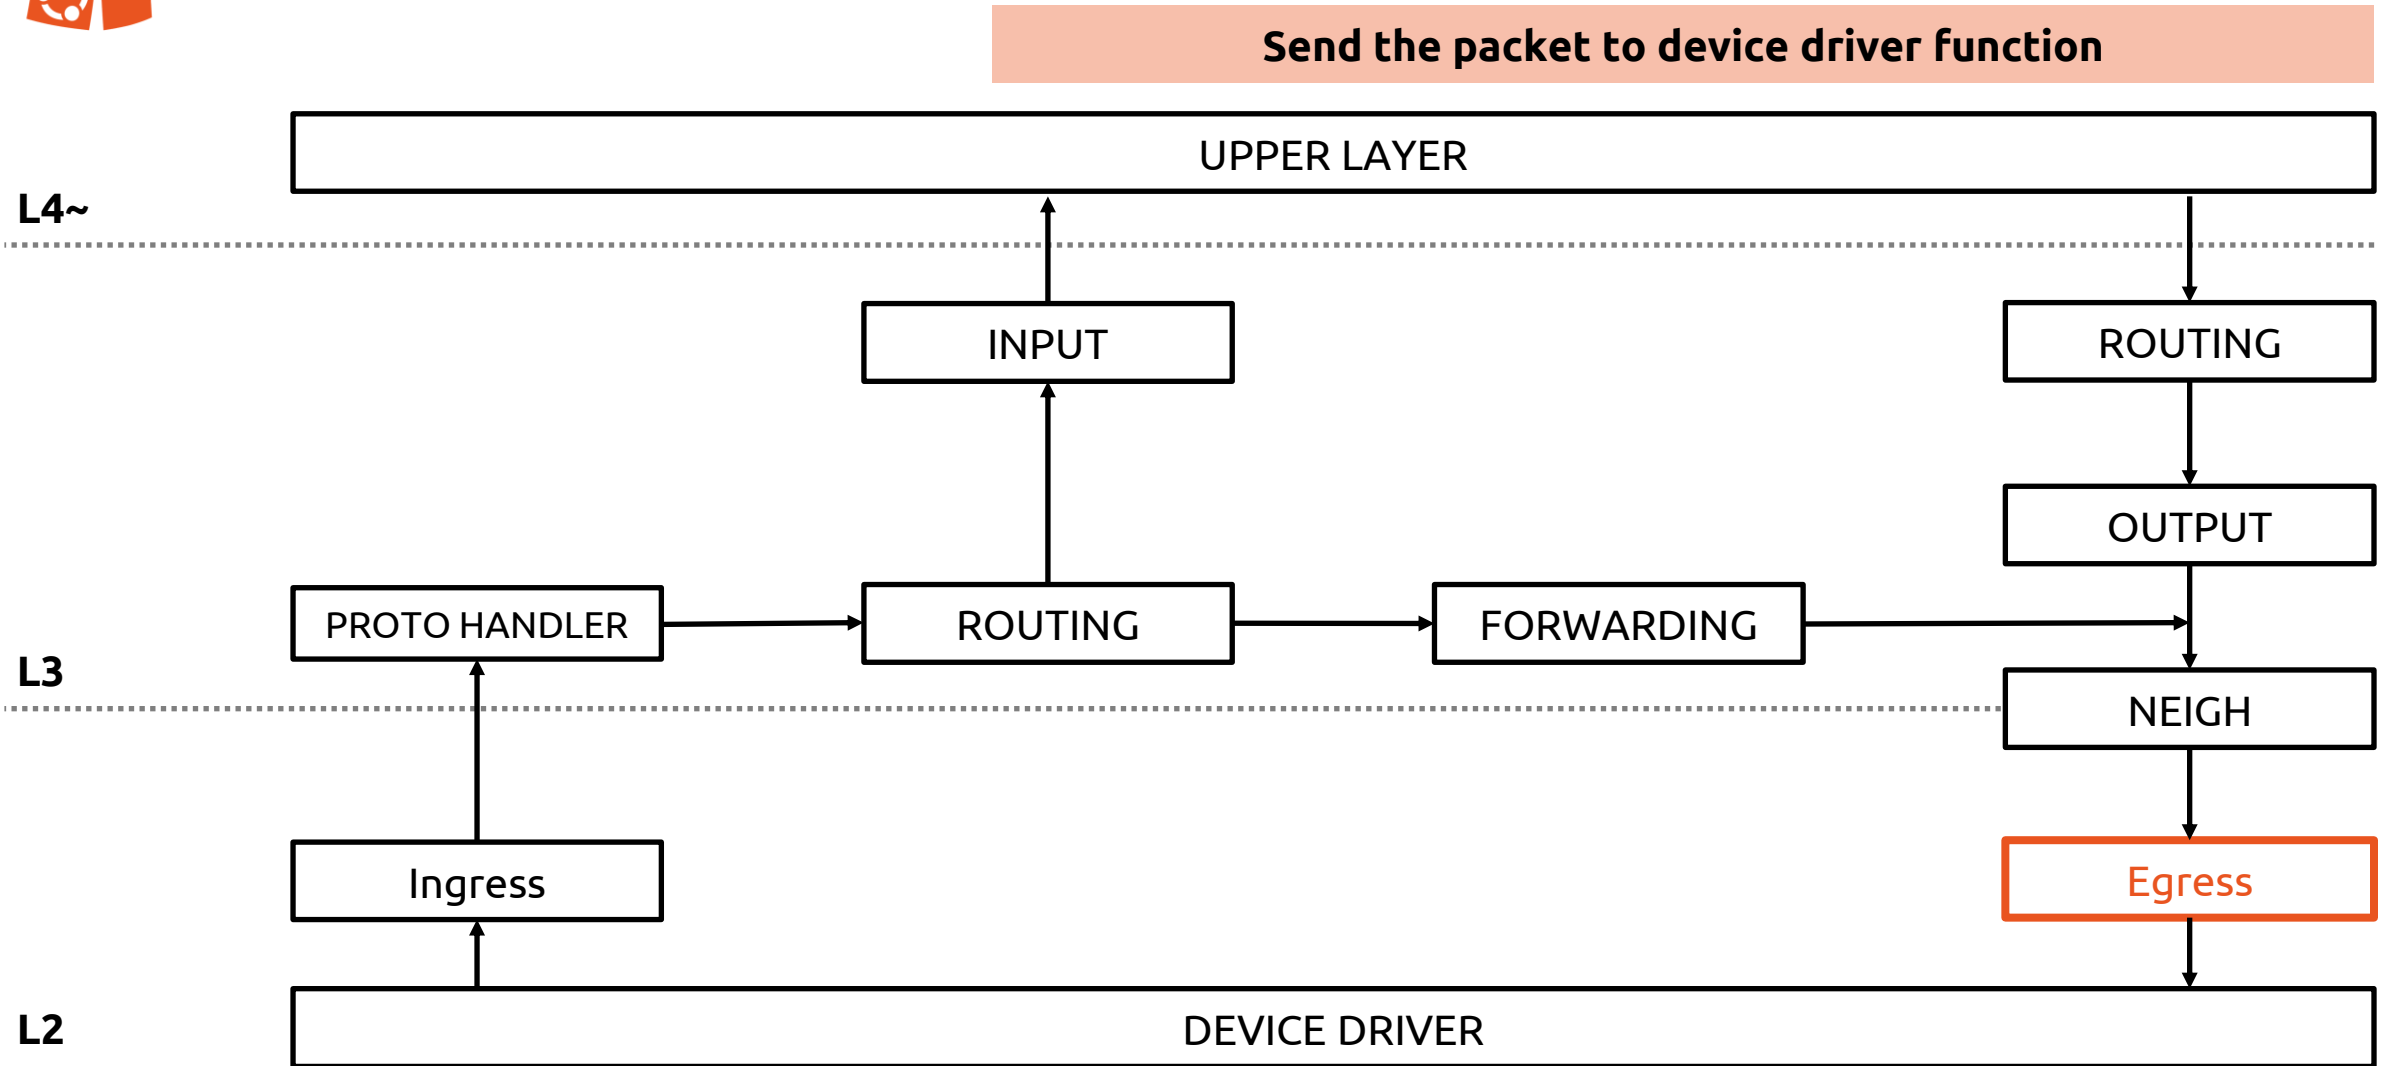


Packet Sending Path



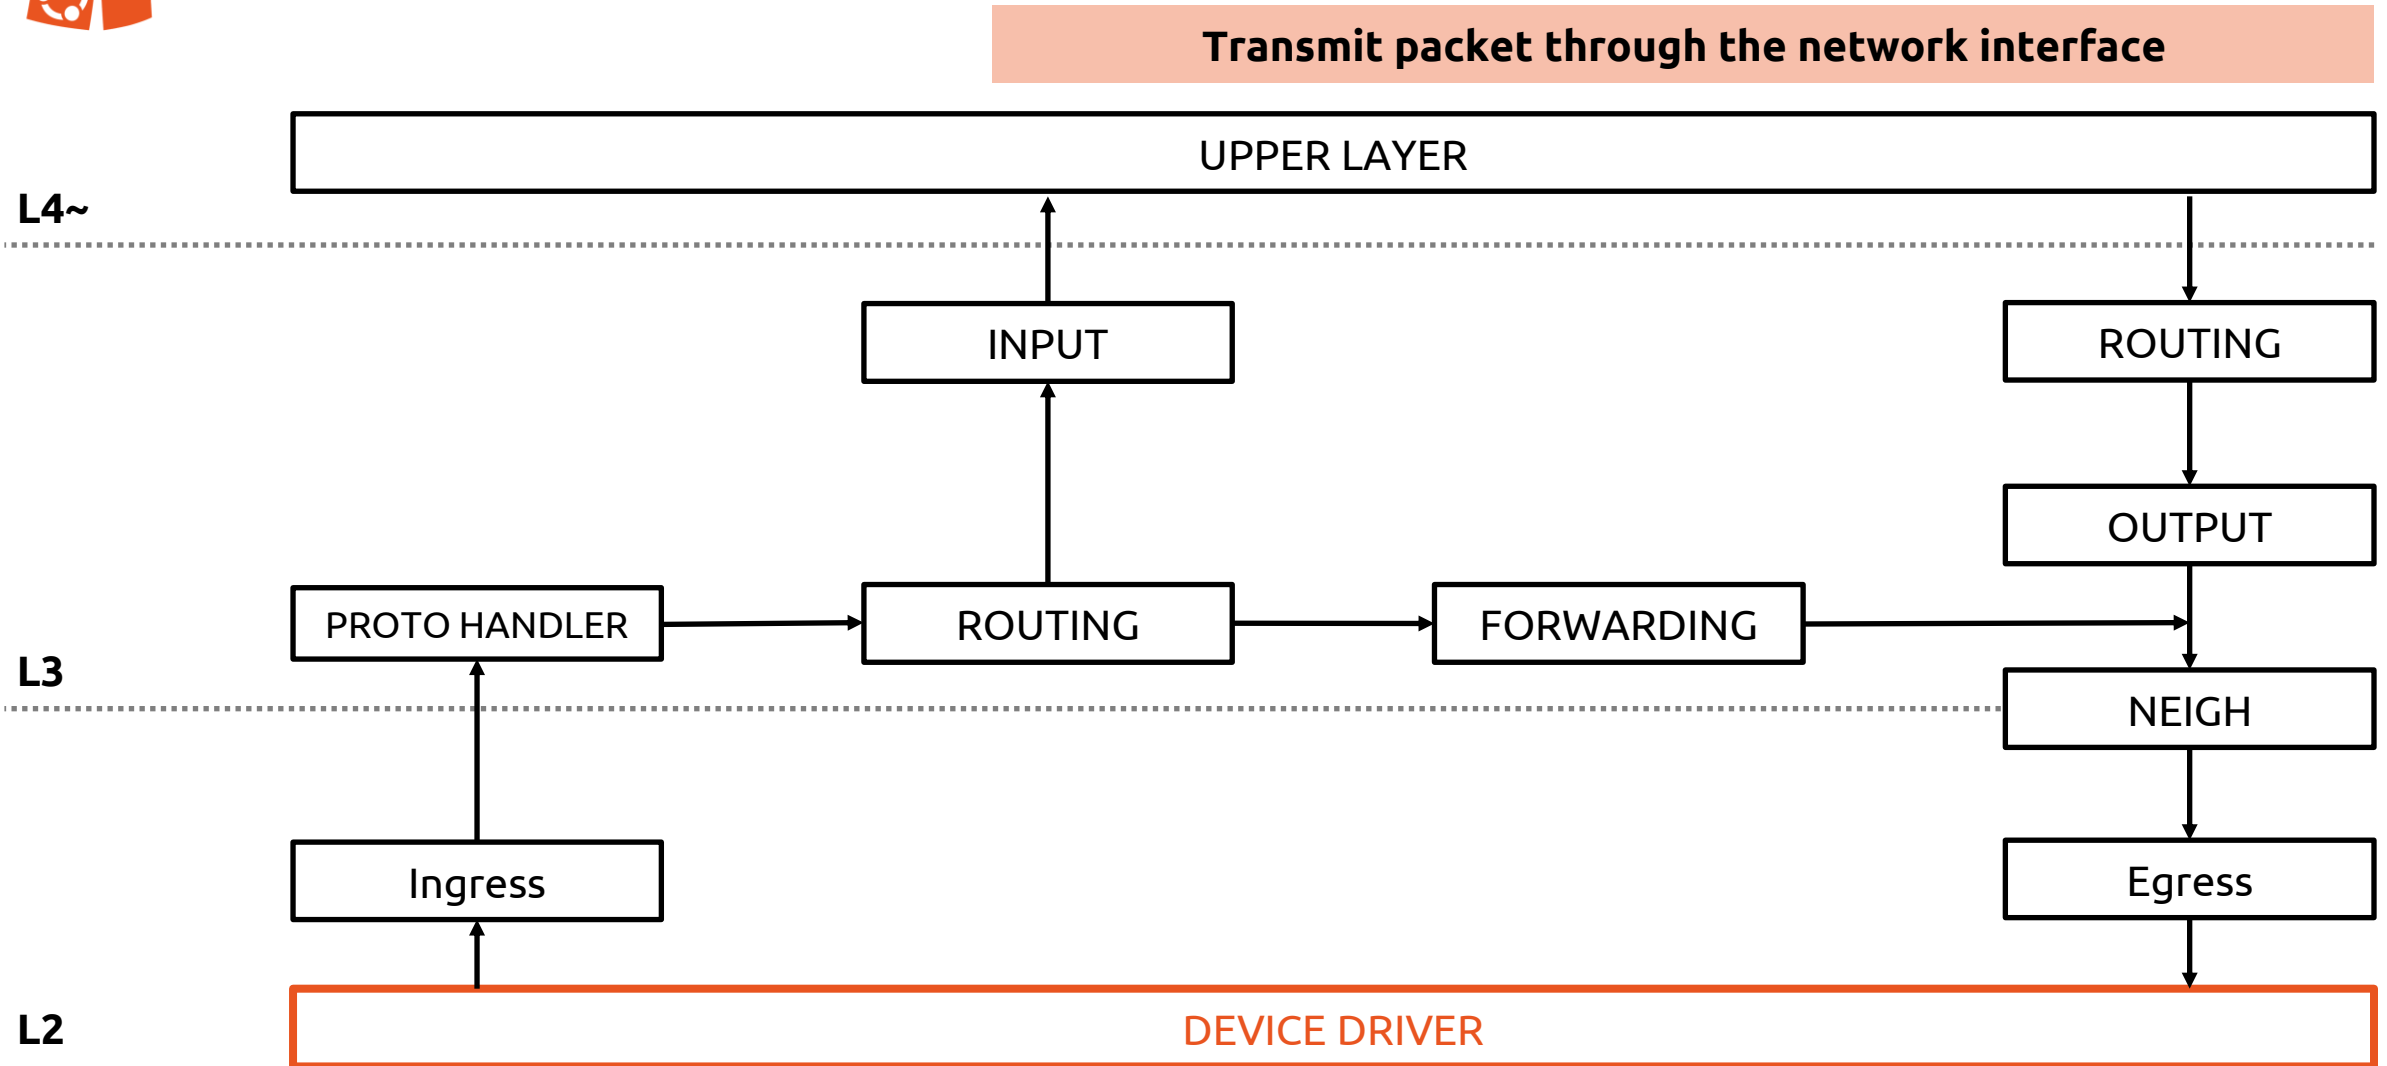


Packet Sending Path



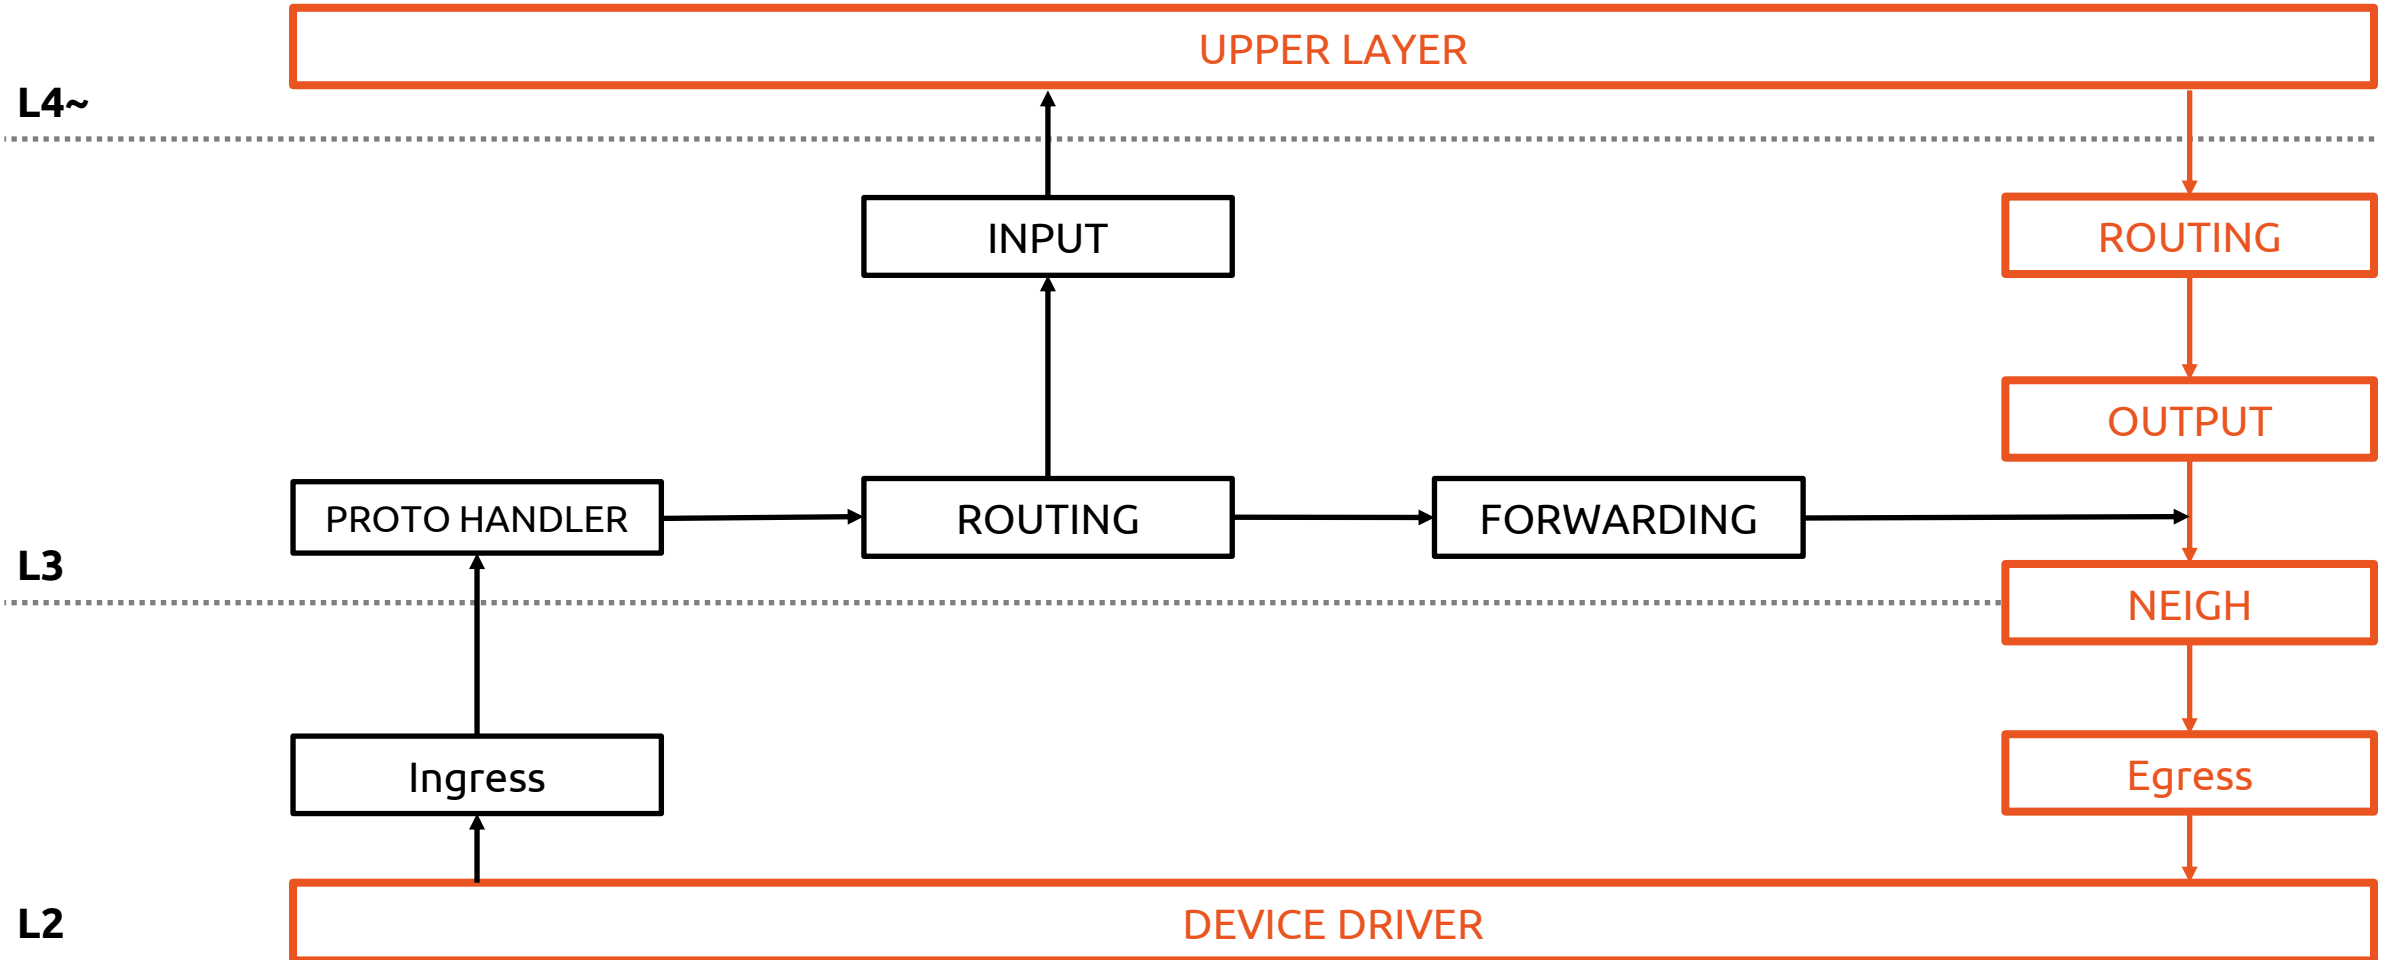


Packet Sending Path






Packet Sending Path






Packet Path

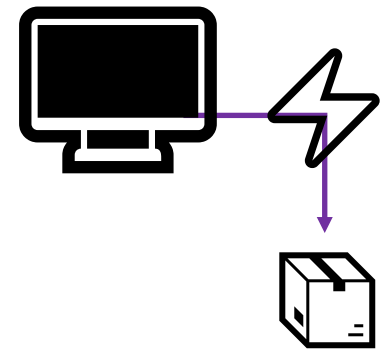
① Receiving Path



② Forwarding Path

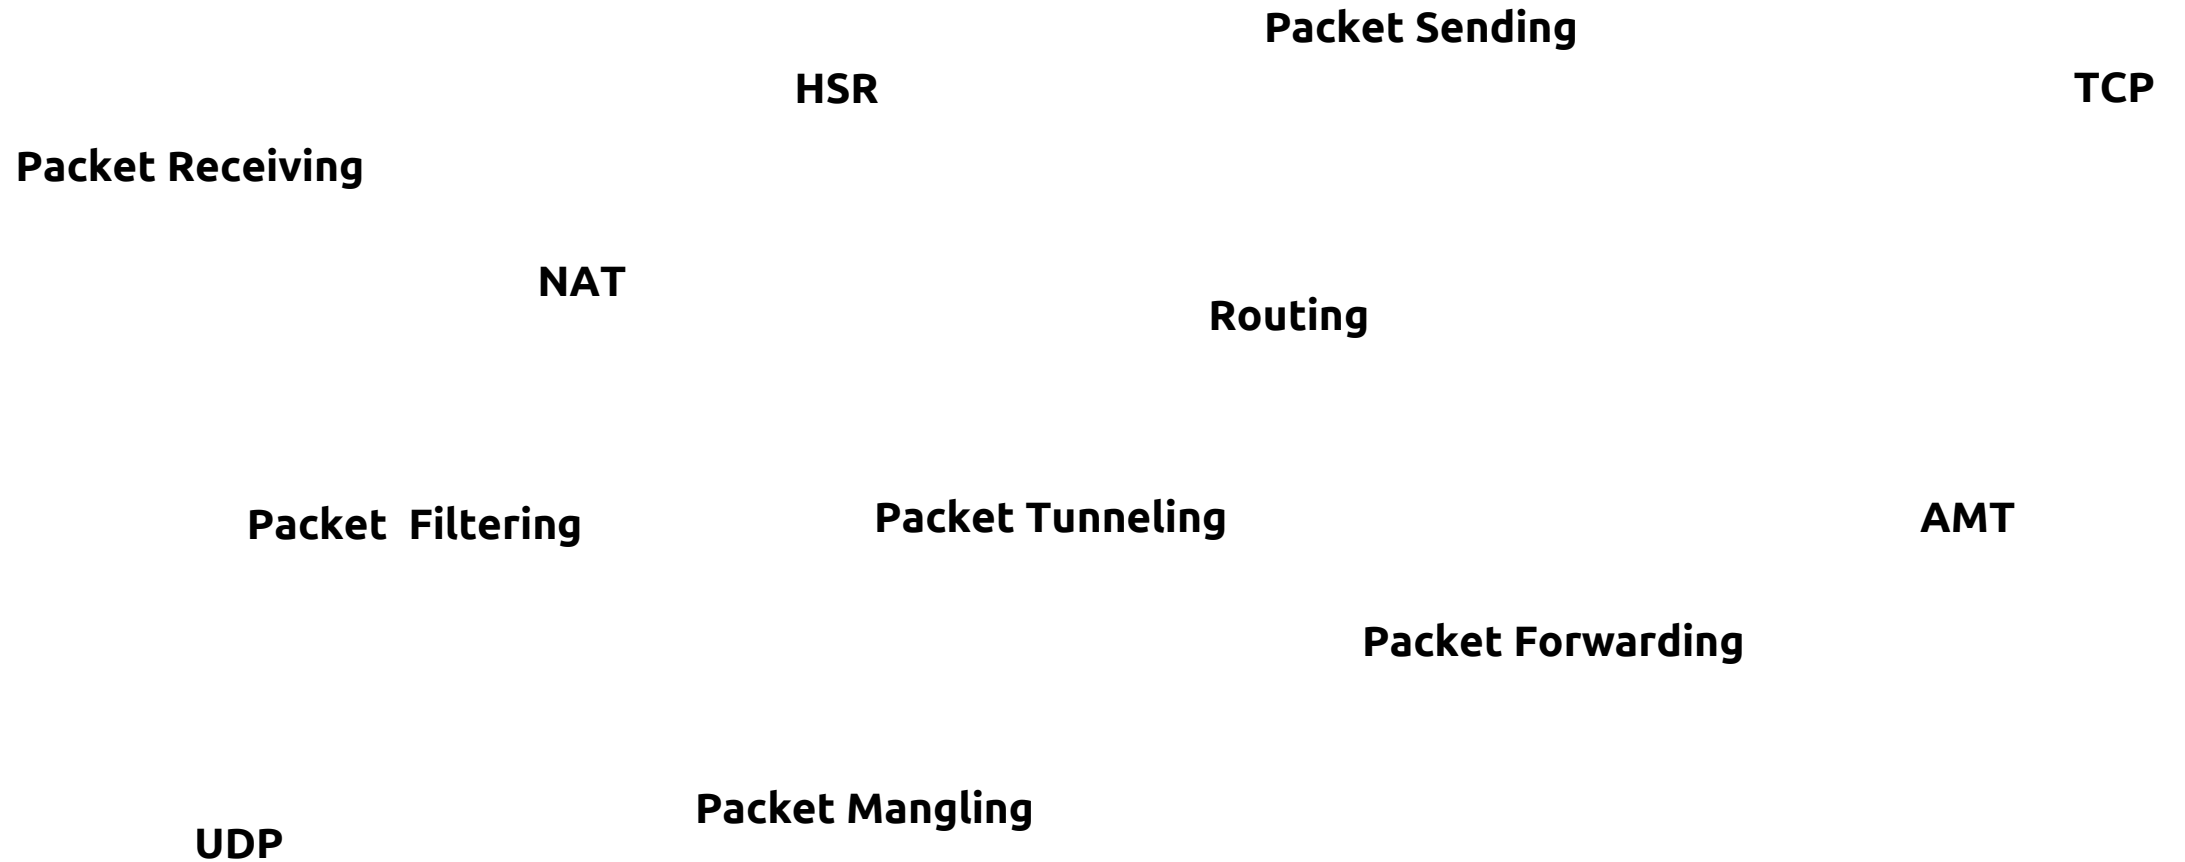


③ Sending Path



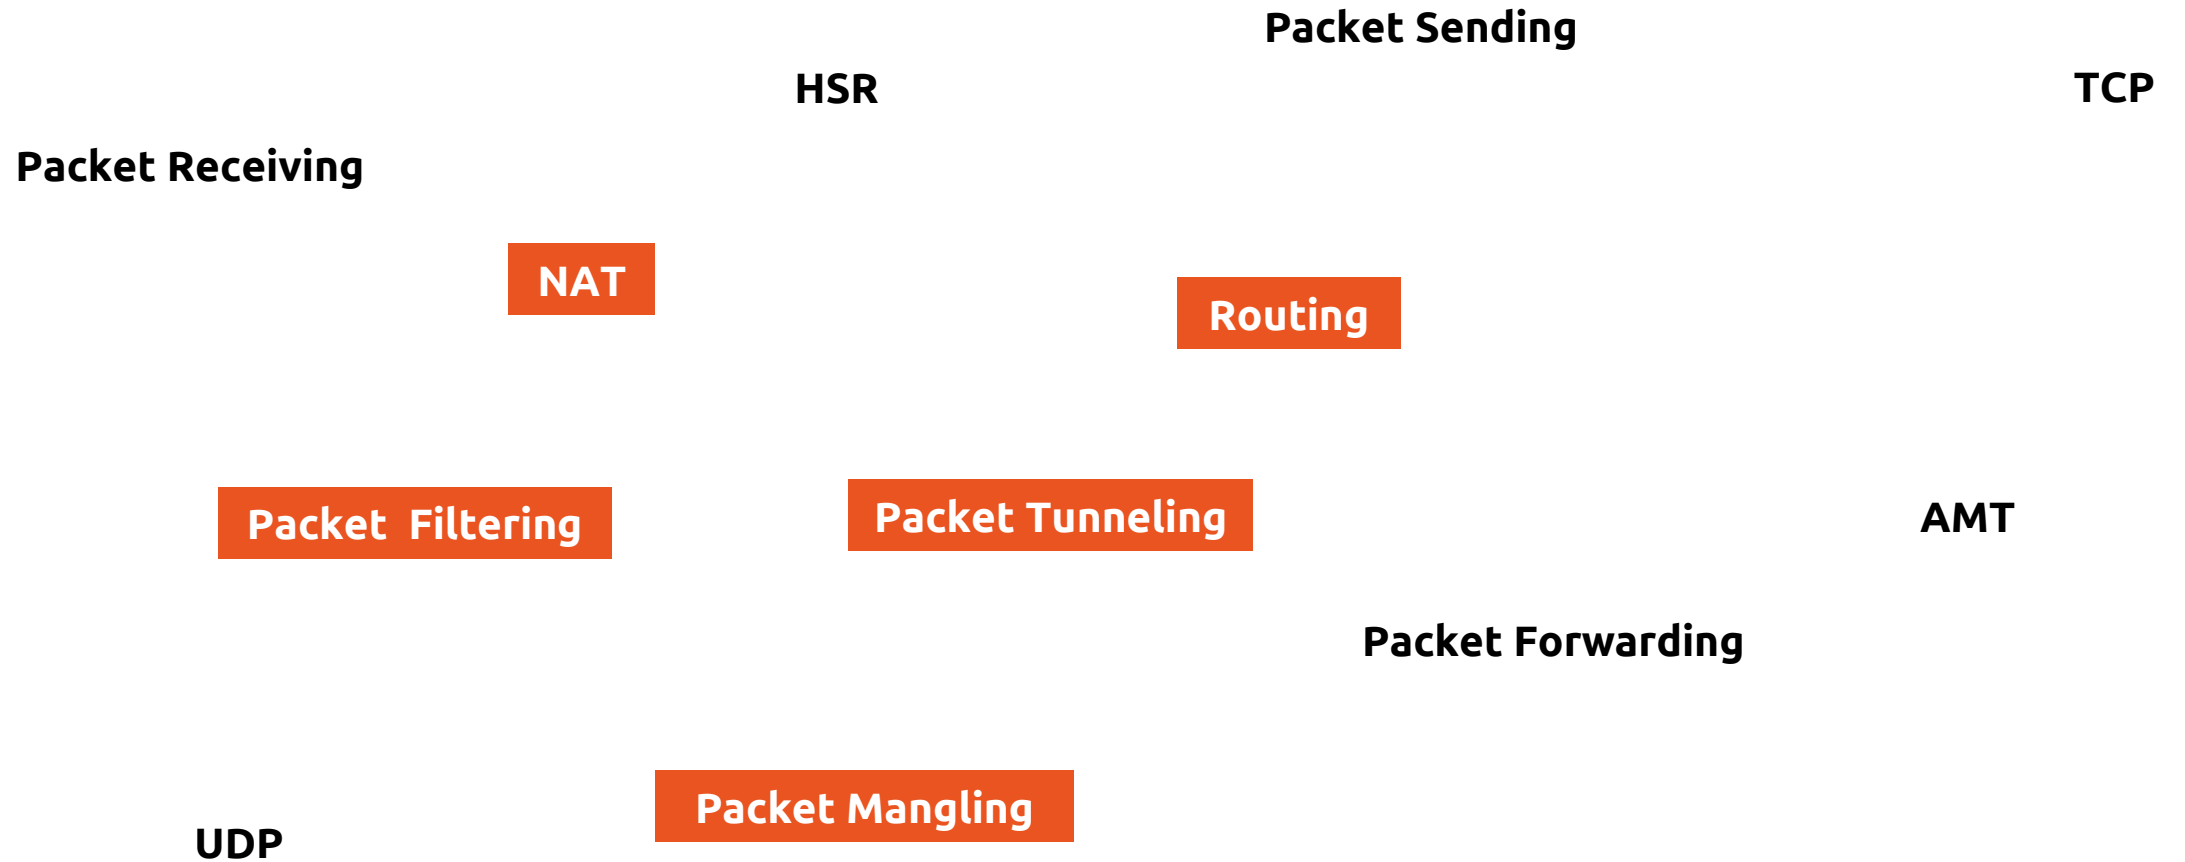


Linux Kernel Networking Stack



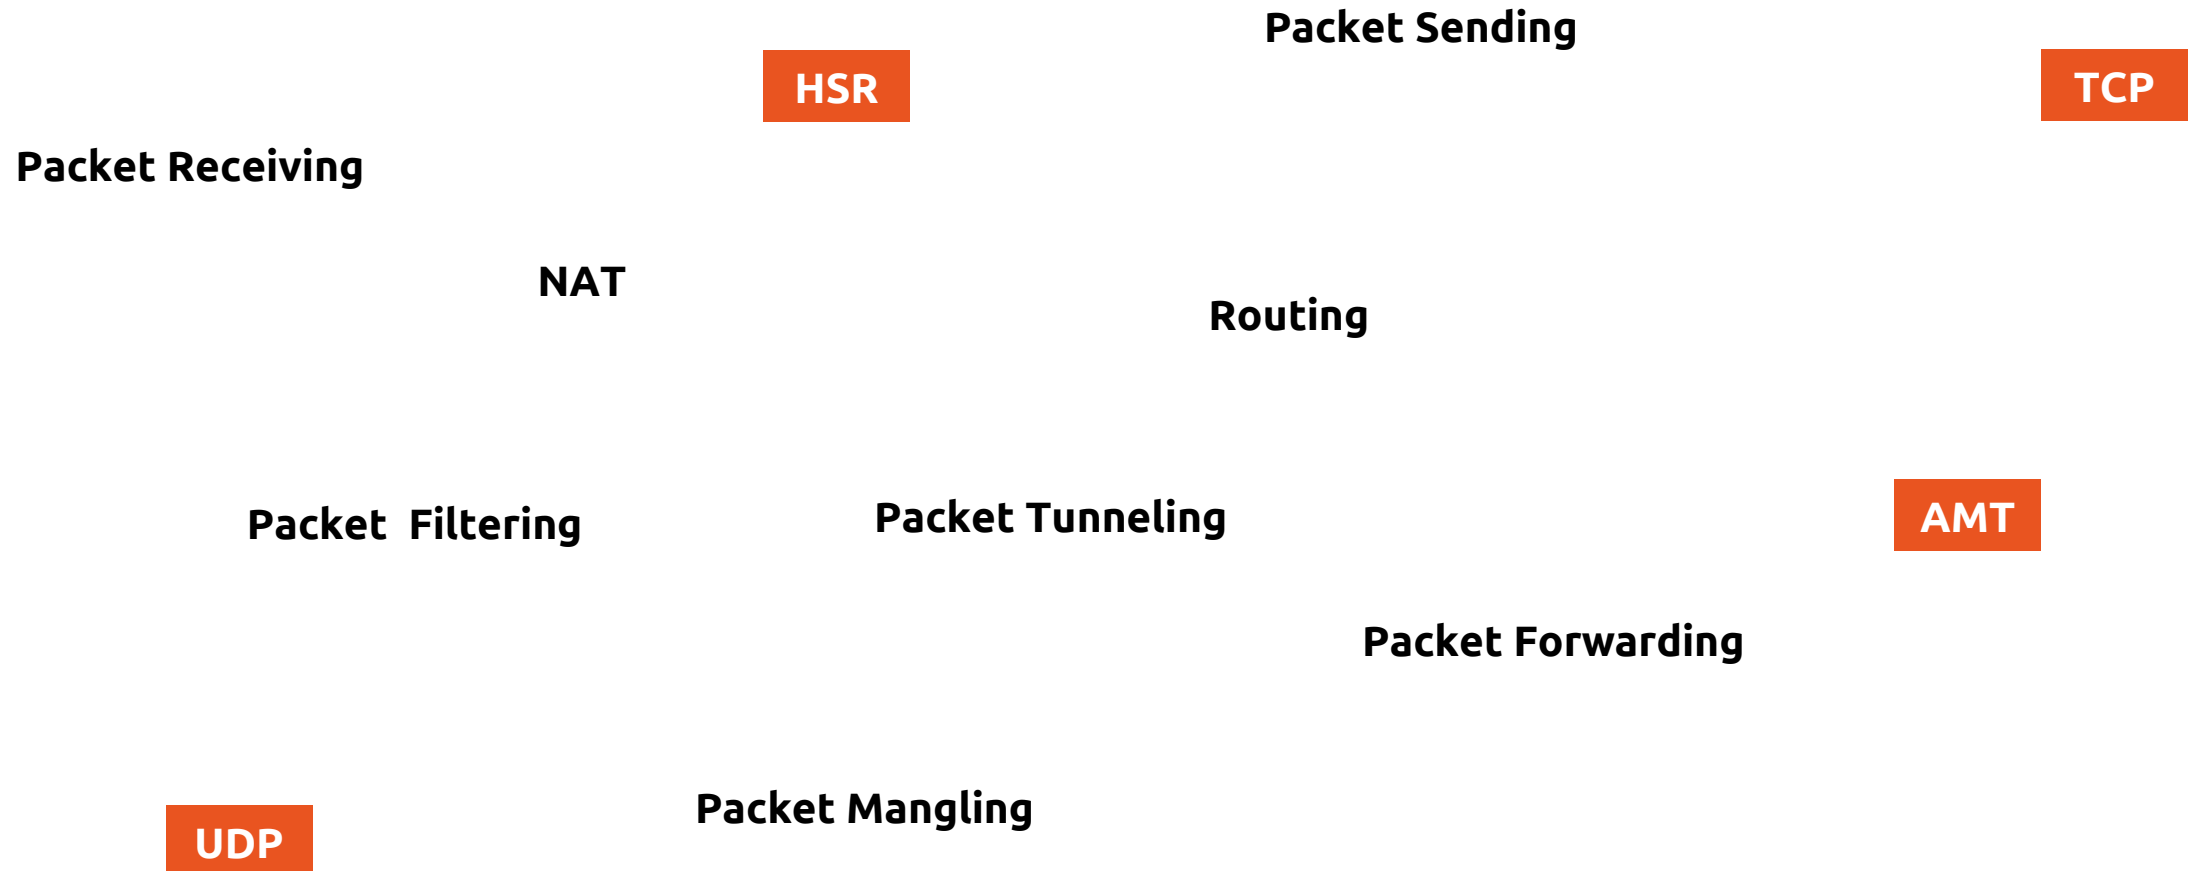


Packet Processing





Network Protocol





2


Kernel Ecosystem



Linux Kernel

The core of the Linux Operating System

Interface that connects computer hardware to processes



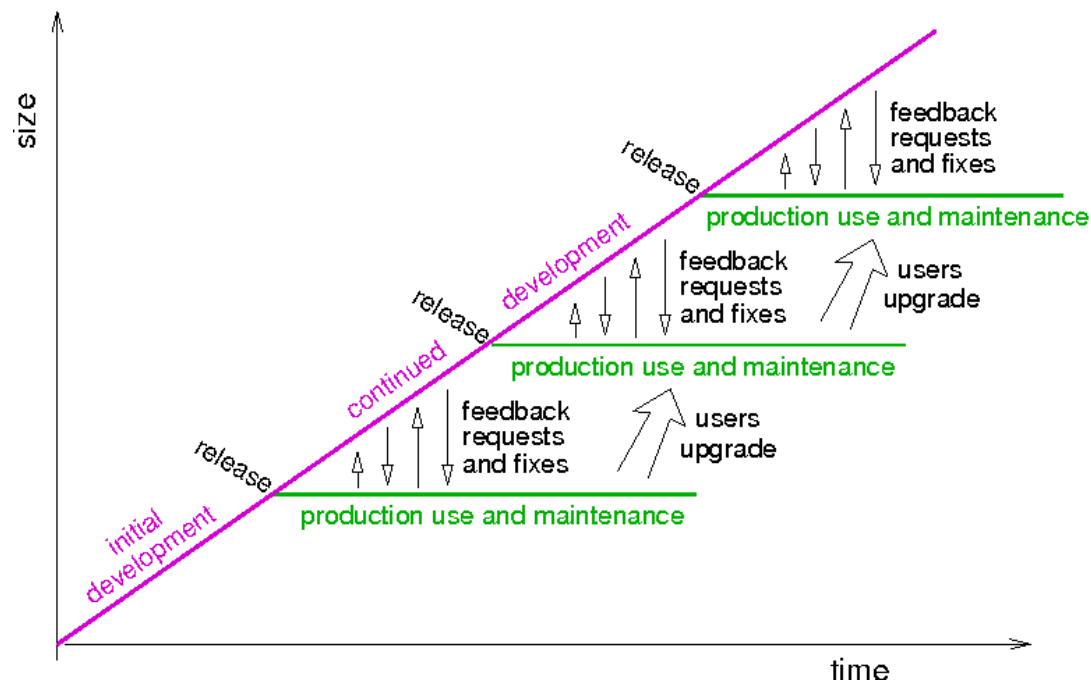
Linux Kernel mascot Tux



Linux Kernel Development

The largest open-source project involving developers worldwide


Development is done in units called **Patch**



<https://www.semanticscholar.org/paper/Perpetual-development%3A-A-model-of-the-Linux-kernel-Feitelson/eb7cf446e983e98c1400c8181949f038caf0c8a8/figure/1>



Kernel Release



new kernel released in every 10-11 weeks (2 months), stable kernel released once a week

Mainline Kernel

Stable

Long-term

Release Candidate (RC)



Kernel Release - 1. Mainline Kernel



- released every 9 to 10 weeks
- **New features introduced, major bug fix patches**



Kernel Release - 2. Stable



Stable

- Switch to stable mode after mainline kernel release
- **Bug fixes** & mainline kernel backporting



Kernel Release - 3. LTS Release

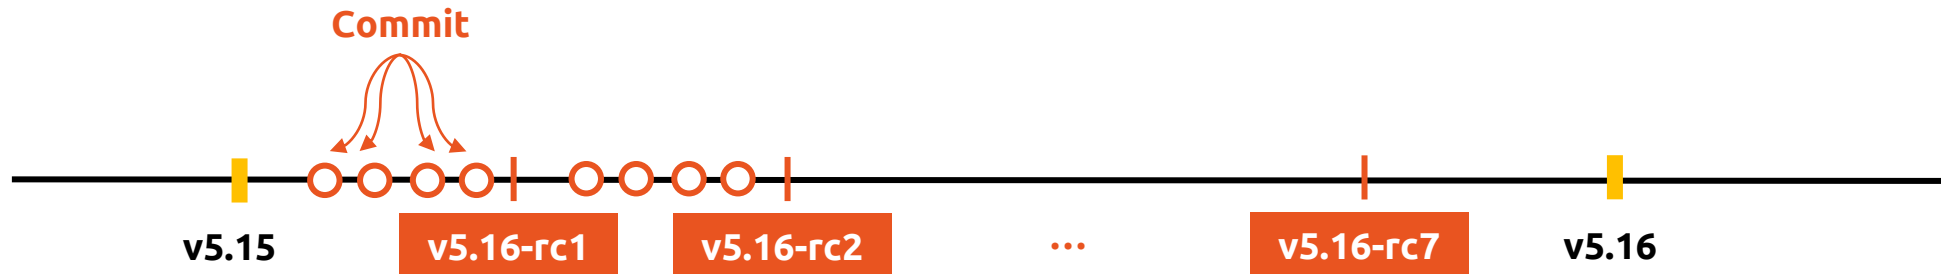


LTS Release

- provides a **critical bug fix** among the stable release versions



Kernel Release - 4. Release Candidate (RC)



Release Candidate (RC)

- "**pre-release**" of Mainline kernel
- For **testing new features**
- This release **must be compiled** from source
- Usually, released up to rc7




Linux Kernel Release Cycle



5.15 mainline
kernel version release



Linux Kernel Release Cycle




5.15 mainline
kernel version release

5.15 Stable release



Linux Kernel Release Cycle




5.15 mainline
kernel version release

5.15 Stable release

5.15 LTS release



Linux Kernel Release Cycle




5.15 mainline kernel version release

5.15 Stable release

5.15 LTS release

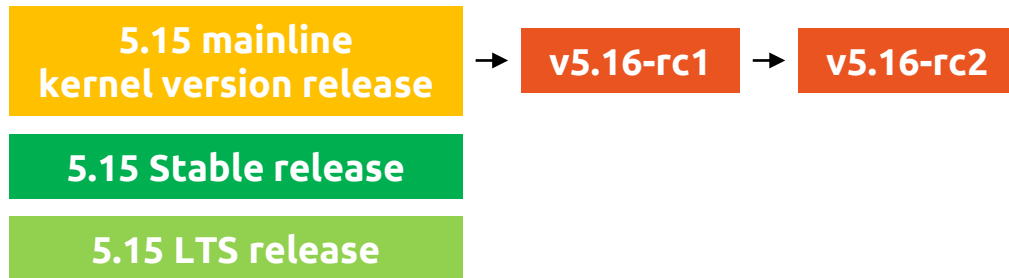



Linux Kernel Release Cycle






Linux Kernel Release Cycle







Linux Kernel Release Cycle



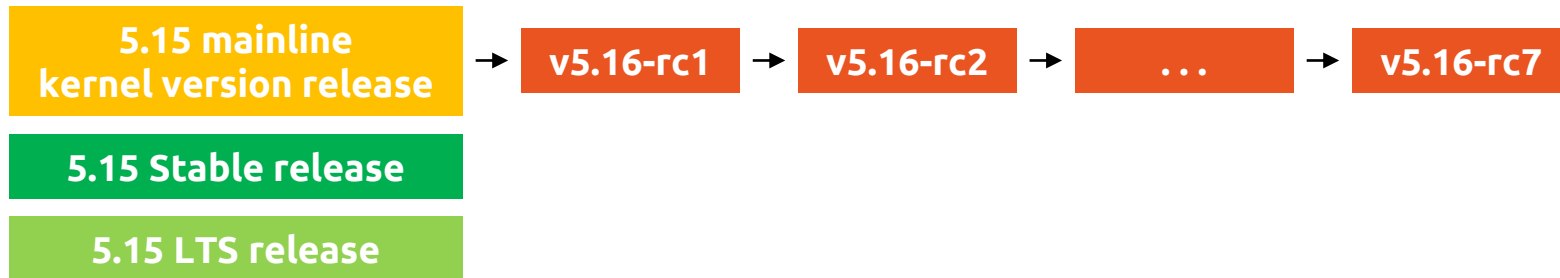



Linux Kernel Release Cycle





Linux Kernel Release Cycle







Linux Kernel Release Cycle



<http://vger.kernel.org/~davem/net-next.html>




Linux Kernel Release Cycle





Linux Kernel Release Cycle





*** Merge window (2 weeks)**

- Linux Torvalds receives a pull of code from the maintainer of subsystem and adds it together
- New version of the kernel is released after merge window

- After period ends, the kernel is developed every two weeks.
Naming convention: <mainline kernel version name>-rc-<rc version name>



Linux Kernel Release Cycle





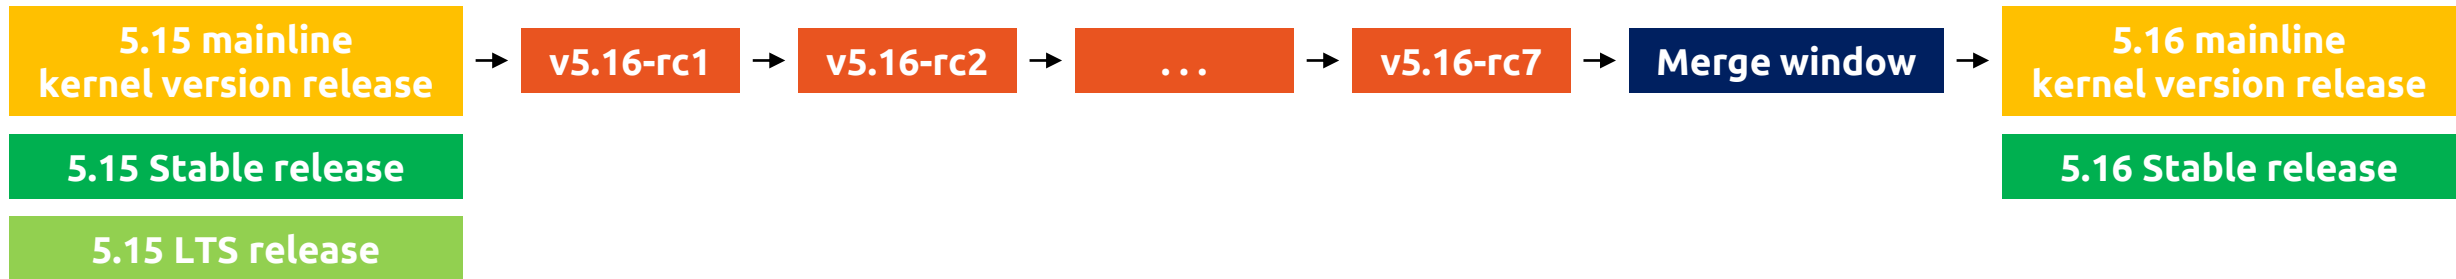

Linux Kernel Release Cycle



<http://vger.kernel.org/~davem/net-next.html>





Linux Kernel Release Cycle





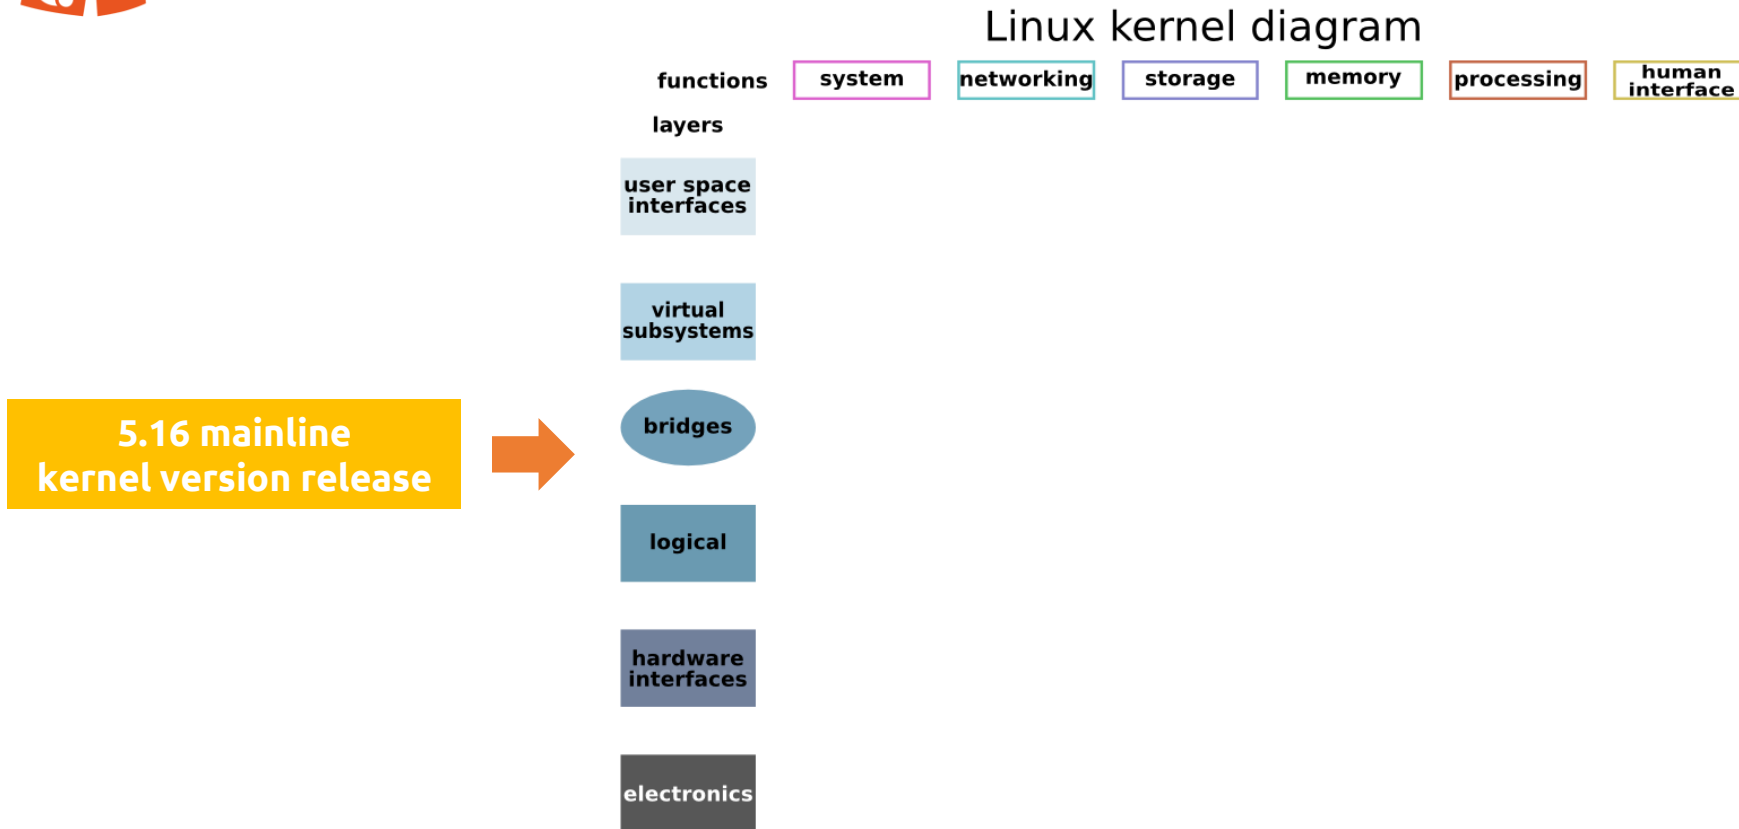
Linux Kernel Release Cycle





Linux Kernel Subsystem

<http://www.MakeLinux.net/kernel/diagram>

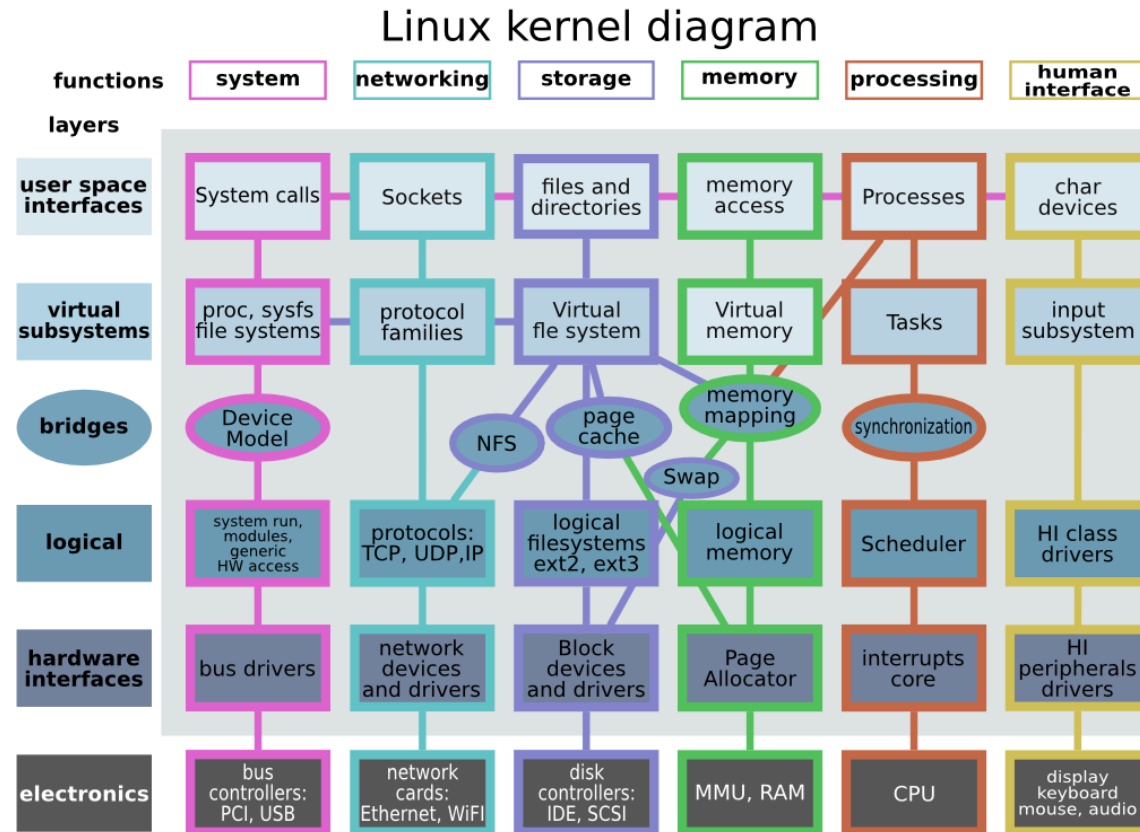


Inside the kernel there are **numerous subsystems,**



Linux Kernel Subsystem

<http://www.MakeLinux.net/kernel/diagram>



5.16 mainline kernel version release




Inside the kernel there are numerous subsystems, Features developed by each subsystem are combined to be released as the mainline kernel



Linux Kernel Subsystem

<http://www.MakeLinux.net/kernel/diagram>



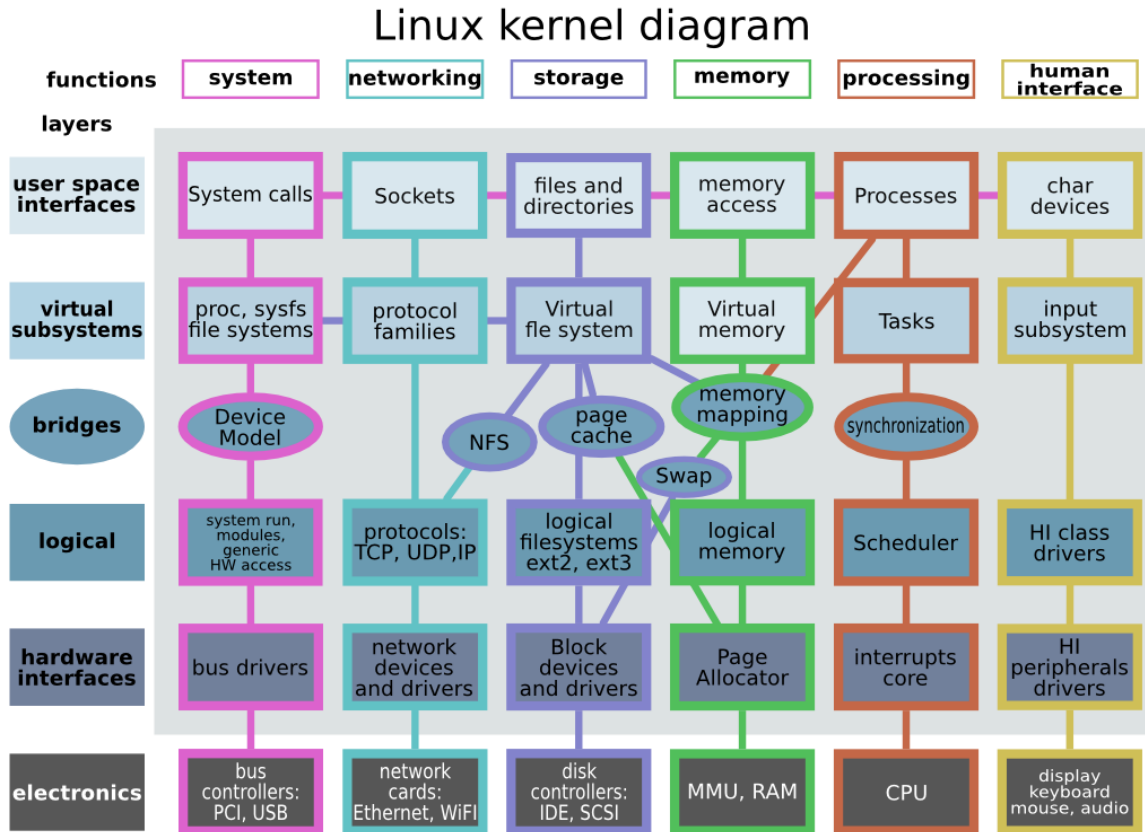
```
linux $ tree -d -L 1
├── arch
├── block
├── certs
├── crypto
├── Documentation
├── drivers
├── fs
├── include
├── init
├── ipc
├── kernel
├── lib
├── LICENSES
├── mm
├── net
├── samples
├── scripts
├── security
├── sound
├── tools
├── usr
└── virt
```

Subsystems are developed in **each directory** of the source code.



Linux Kernel Subsystem

<http://www.MakeLinux.net/kernel/diagram>



```
linux $ tree -d -L 1
```

```
├── arch
├── block           => block device subsystem
├── certs
├── crypto
├── Documentation
├── drivers         => driver subsystem
├── fs             => file system subsystem
├── include
├── init
├── ipc           => IPC subsystem
├── kernel
├── lib
├── LICENSES
├── mm           => memory management subsystem
├── net          => network subsystem
├── samples
├── scripts
├── security
├── sound
├── tools
├── usr
└── virt
```

Network subsystem are developed in net directory of the source code.



Kernel Tree



Kernel.org git repositories

Git repositories hosted at kernel.org (singapore)

<https://git.kernel.org/>

index

search

Name	Description	Owner	Idle	Links
<i>pub/scm/bluetooth</i>				
bluetooth-next.git	Unnamed repository; edit this file 'description' to name the repository.	Grokmirror user	19 months	summary log tree
bluez-hcidump.git	Bluetooth packet analyzer	Marcel Holtmann	5 years	summary log tree
bluez.git	Bluetooth protocol stack for Linux	Marcel Holtmann	20 hours	summary log tree
obexd.git	OBEX Server	Marcel Holtmann	5 years	summary log tree
sbcc.git	Bluetooth low-complexity, subband codec (SBC) library	Marcel Holtmann	4 months	summary log tree
<i>pub/scm/boot</i>				
dracut/dracut.git	dracut - Initramfs generator using udev	Harald Hoyer	8 months	summary log tree
efilinux/efilinux.git	The efilinux UEFI boot loader	Matt Fleming	8 years	summary log tree
syslinux/syslinux.git	Unnamed repository; edit this file 'description' to name the repository.	Syslinux workgroup	6 years	summary log tree
<i>pub/scm/devel</i>				
pahole/pahole.git	Pahole and other DWARF utils	Arnaldo Carvalho de Melo	21 hours	summary log tree
sparse/chris/sparse.git	Chris Li's sparse repository.	Christopher Li	5 years	summary log tree
sparse/sparse-dev.git	Sparse's development tree	Luc Van Oostenryck	4 months	summary log tree
sparse/sparse-logs.git	Raw logs of warnings of Sparse running on the kernel with summary logs.	Luc Van Oostenryck	4 months	summary log tree
sparse/sparse.git	C semantic parser	Luc Van Oostenryck	4 months	summary log tree
<i>pub/scm/docs</i>				
docsko/korg.git	Source of korg.docs.kernel.org	Doc Group	8 months	summary log tree
kernel/kernel-docs.git	Kernel Documentation tree	Doc Group	9 years	summary log tree
kernel/ksmap.git	Kernel.org keysign map source	Kernel.org users	3 weeks	summary log tree
kernel/pgpkeys.git	Kernel developers PGP keys	Kernel.org Admin	3 days	summary log tree

**kernel code is managed by several major subsystem git repositories,
which is called the **Kernel Tree****



Kernel Tree

Mainline Kernel tree

- Managed by [Linus Torvalds](#)
- Mainline kernel tree & RC release source code

Kernel
Mainline



* <https://www.devcurmudgeon.com/images/mainline-lts-ltsi-genivi-20131025.pdf#page=18>



Kernel Tree

Mainline Kernel tree


- Managed by Linus Torvalds
- Mainline kernel tree & RC release source code

Stable Kernel tree

- Managed by [Greg Kroah-Hartman](#)
- stable release branches source code

Kernel
Mainline

Kernel.org
(Greg K-H)



* <https://www.devcurmudgeon.com/images/mainline-lts-ltsi-genivi-20131025.pdf#page=18>



Kernel Tree

Mainline Kernel tree

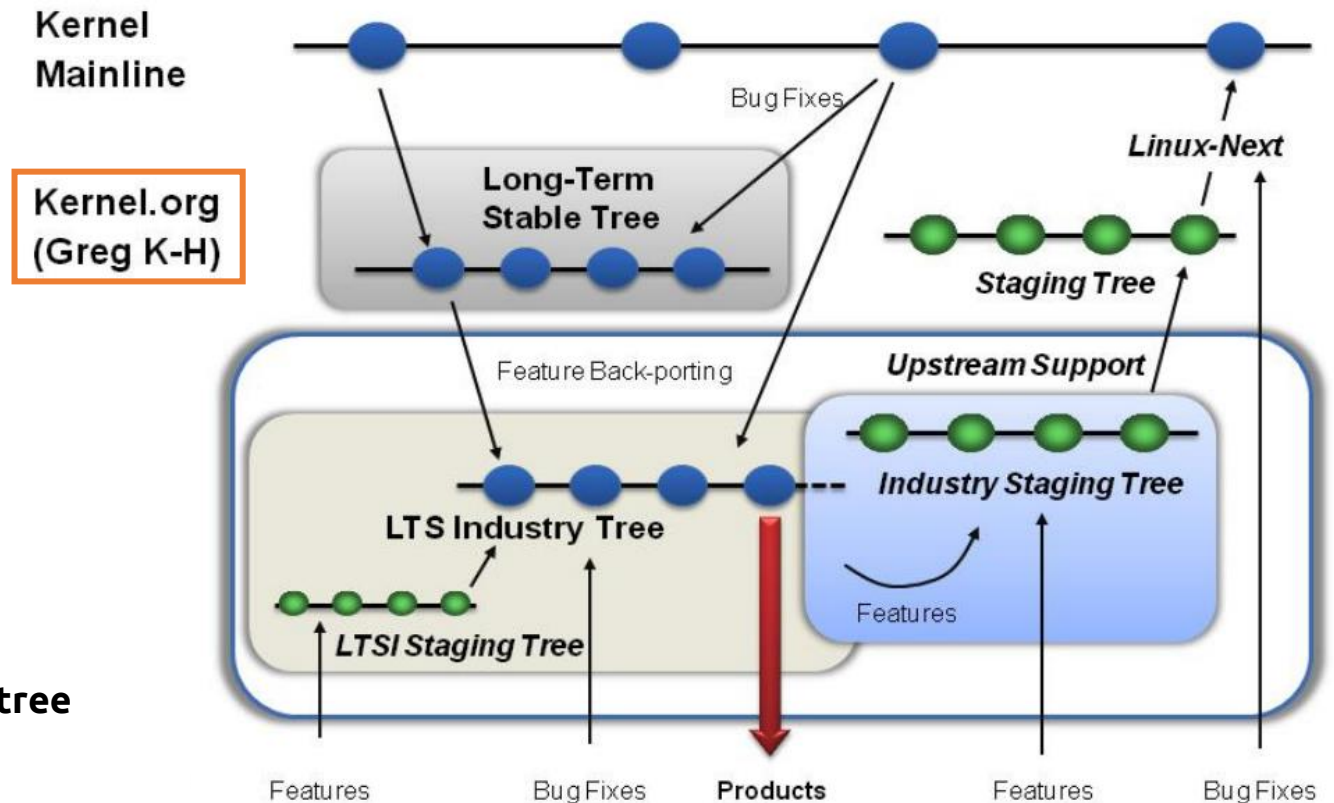
- Managed by Linus Torvalds
- Mainline kernel tree & RC release source code

Stable Kernel tree

- Managed by Greg Kroah-Hartman
- stable release branches source code

linux-next tree

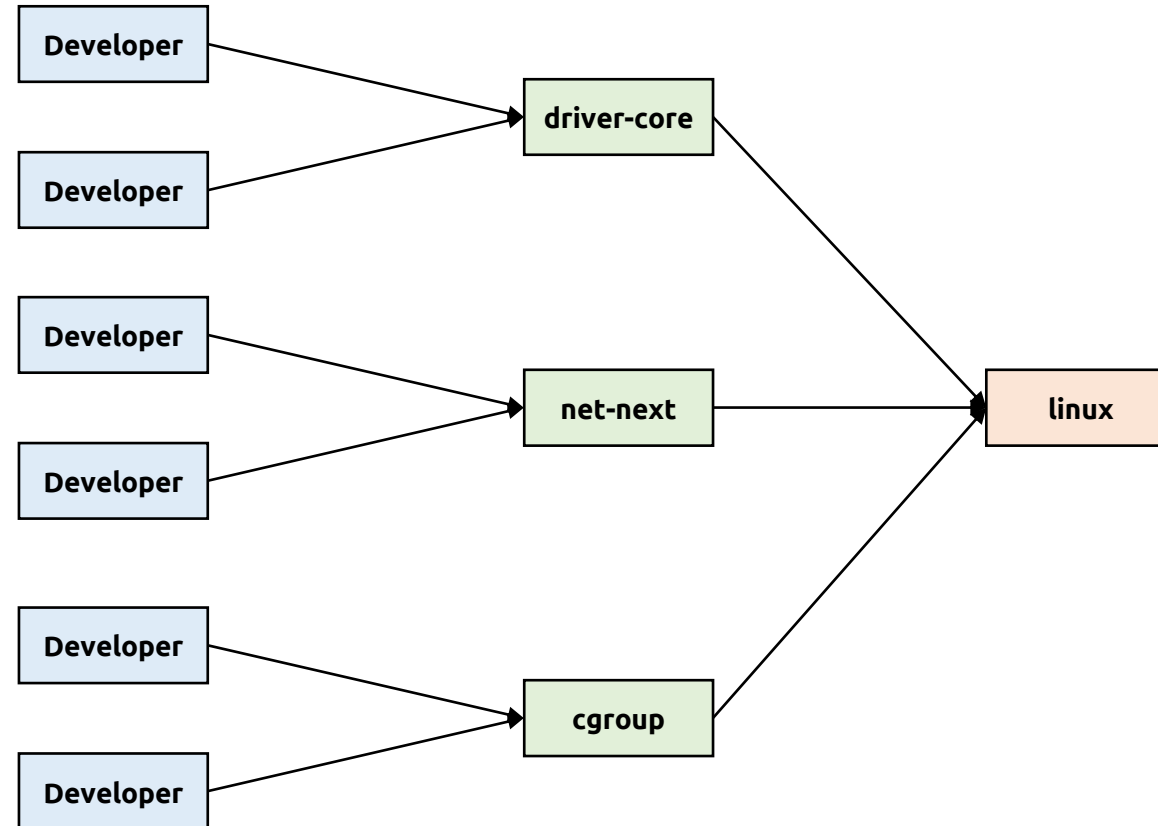
- Managed by [Stephen Rothwell](#)
- Subsystem changes are periodically reflected
- Integrated testing of the code is conducted
- After testing, pull request is performed with mainline kernel tree



* <https://www.devcurmudgeon.com/images/mainline-lts-ltsi-genivi-20131025.pdf#page=18>



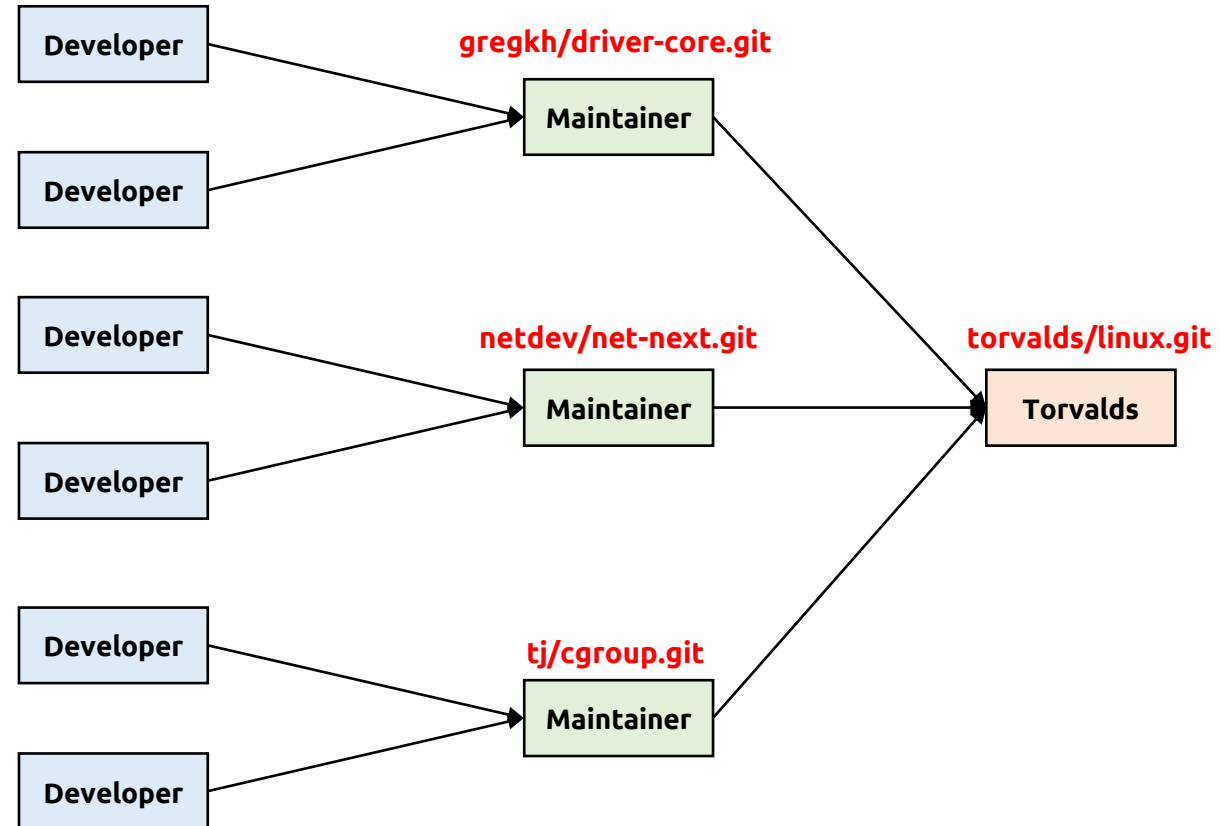
Linux Kernel Tree Hierarchy



Major subsystems have **independent tree**



Linux Kernel Tree Hierarchy




Each subsystem tree is managed by "**Maintainer**"



Linux Kernel Tree Hierarchy

Networking Subsystem



These maintainers **manage the tree** for each subsystem,
and there are **two trees** in each subsystem



Networking Subsystem Tree



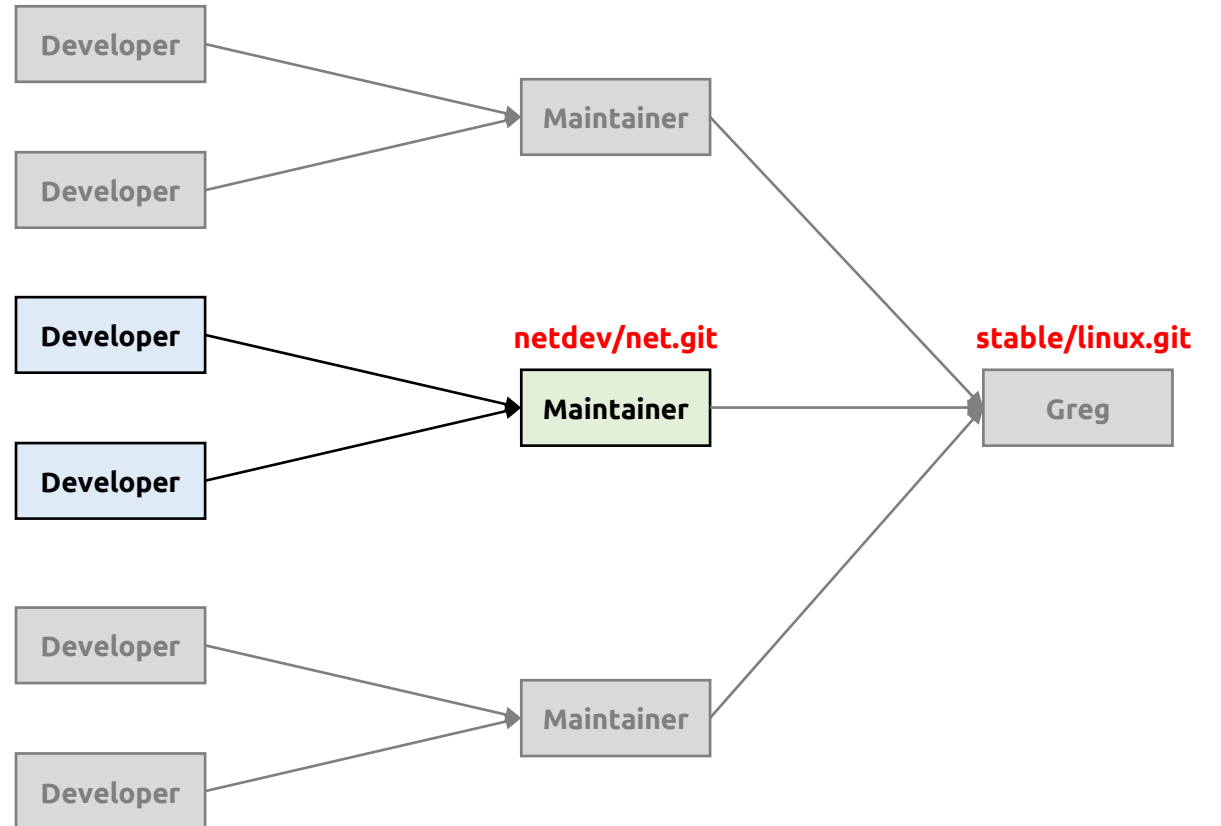
index : kernel/git/netdev/net.git

git.kernel.org/pub/scm/linux/kernel/git/netdev/net.git

Netdev Group

netdev/net.git : netdev Stable Tree

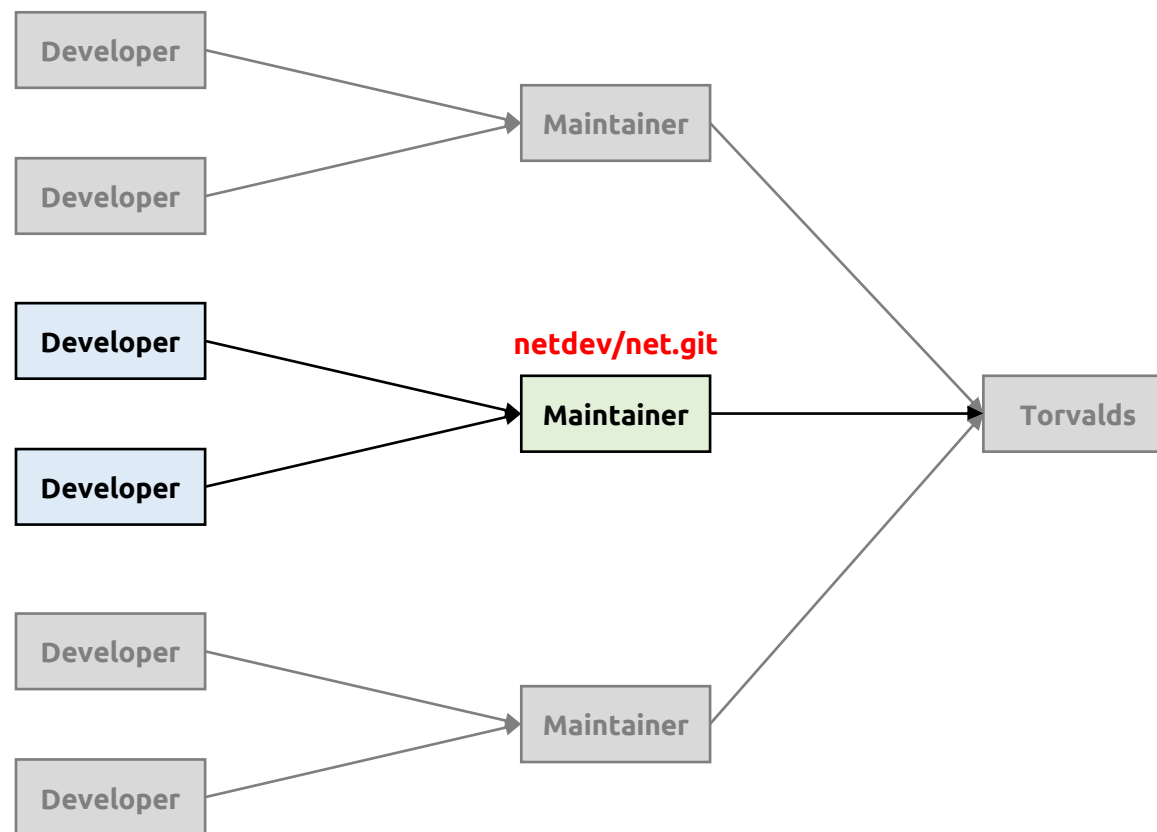
Allow only release version of bug fix patch





Subsystem Tree

```
$ cat MAINTAINERS | grep -A 30 'NETWORKING \[GENERAL]'  
NETWORKING [GENERAL]  
M: "David S. Miller" <davem@davemloft.net>  
M: Eric Dumazet <edumazet@google.com>  
M: Jakub Kicinski <kuba@kernel.org>  
M: Paolo Abeni <pabeni@redhat.com>  
L: netdev@vger.kernel.org  
S: Maintained  
Q: https://patchwork.kernel.org/project/netdevbpf/list/  
B: mailto:netdev@vger.kernel.org  
T: git git://git.kernel.org/.../netdev/net.git  
T: git git://git.kernel.org/.../netdev/net-next.git  
F: Documentation/networking/  
F: Documentation/process/maintainer-netdev.rst  
F: include/linux/in.h  
F: include/linux/net.h  
F: include/linux/netdevice.h  
F: include/net/  
F: include/uapi/linux/in.h  
F: include/uapi/linux/net.h  
F: include/uapi/linux/net_namespace.h  
F: include/uapi/linux/netdevice.h  
F: lib/net_utils.c  
F: lib/random32.c  
F: net/  
F: tools/testing/selftests/net/
```





Mailing List

Each subsystem has a **mailing list**, where contributors submit and **discuss patches**.

<http://vger.kernel.org/vger-lists.html>

```
$ cat MAINTAINERS | grep -A 30 'NETWORKING \[GENERAL]'
NETWORKING [GENERAL]
M: "David S. Miller" <davem@davemloft.net>
M: Eric Dumazet <edumazet@google.com>
M: Jakub Kicinski <kuba@kernel.org>
M: Paolo Abeni <pabeni@redhat.com>
L: netdev@vger.kernel.org
S: Maintained
```

Majordomo lists at VGER.KERNEL.ORG

REMEMBER: Subscription to these lists go via <majordomo@vger.kernel.org> !

Note about archives: Listed archives are those that have been reported to vger's maintainers, or that we have found out otherwise. As things are, *list of archives is not complete*.

[alsa-devel](#), [autofs](#), [backports](#), [bpf](#), [ceph-devel](#), [cgroups](#), [cpufreq](#), [dash](#), [dccp](#), [devicetree-compiler](#), [devicetree-spec](#), [devicetree](#), [dmaengine](#), [dwarves](#), [ecryptfs](#), [fio](#), [fstests](#), [git-commits-24](#), [git-commits-head](#), [git](#), [hail-devel](#), [initramfs](#), [io-uring](#), [irda-users](#), [kernel-janitors](#), [kernel-packagers](#), [kernel-testers](#), [keyrings](#), [kvm-commits](#), [kvm-ia64](#), [kvm-ppc](#), [kvm](#), [lartc](#), [libzbc](#), [linux-8086](#), [linux-acpi](#), [linux-admin](#), [linux-alpha](#), [linux-api](#), [linux-apps](#), [linux-arch](#), [linux-arm-msm](#), [linux-assembly](#), [linux-bbs](#), [linux-bcache](#), [linux-bcachefs](#), [linux-block](#), [linux-bluetooth](#), [linux-brace](#), [linux-btrfs](#), [linux-c-programming](#), [linux-can](#), [linux-cifs](#), [linux-clk](#), [linux-config](#), [linux-console](#), [linux-coverity](#), [linux-crypto](#), [linux-csky](#), [linux-cxl](#), [linux-diald](#), [linux-doc](#), [linux-edac](#), [linux-efi](#), [linux-embedded](#), [linux-ext4](#), [linux-fbdev](#), [linux-fido](#), [linux-fpga](#), [linux-fscrypt](#), [linux-fsdevel](#), [linux-fsf](#), [linux-ftp](#), [linux-gcc](#), [linux-gpio](#), [linux-hams](#), [linux-hardening](#), [linux-hexagon](#), [linux-hotplug](#), [linux-hwmon](#), [linux-hyperv](#), [linux-i2c](#), [linux-ia64](#), [linux-ibcs2](#), [linux-ide](#), [linux-iiio](#), [linux-input](#), [linux-integrity](#), [linux-ipx](#), [linux-isdn](#), [linux-japanese](#), [linux-kbuild](#), [linux-kernel-announce](#), [linux-kernel-announce-posters](#), [linux-kernel](#), [linux-kseltest](#), [linux-laptop](#), [linux-leds](#), [linux-linuxss](#), [linux-lvm](#), [linux-lugnuts](#), [linux-m68k-cvscommit](#), [linux-m68k](#), [linux-man](#), [linux-mca](#), [linux-media](#), [linux-metag](#), [linux-mips](#), [linux-mmc](#), [linux-modules](#), [linux-msdos-devel](#), [linux-msdos](#), [linux-new-lists](#), [linux-newbie](#), [linux-next](#), [linux-nfs](#), [linux-nilfs](#), [linux-numa](#), [linux-omap](#), [linux-opengl](#), [linux-parisc](#), [linux-pci](#), [linux-perf-users](#), [linux-pm](#), [linux-ppp](#), [linux-pwm](#), [linux-raid](#), [linux-rdma](#), [linux-remoteproc](#), [linux-renesas-soc](#), [linux-rt-users](#), [linux-rtc](#), [linux-s390](#), [linux-samsung-soc](#), [linux-scsi](#), [linux-sctp](#), [linux-security-module](#), [linux-serial](#), [linux-sgx](#), [linux-sh](#), [linux-smp](#), [linux-soc](#), [linux-sound](#), [linux-sparse](#), [linux-spdx](#), [linux-spi](#), [linux-standards](#), [linux-svgalib](#), [linux-tape](#), [linux-tegra](#), [linux-tip-commits](#), [linux-toolchains](#), [linux-trace-devel](#), [linux-trace-kernel](#), [linux-trace-users](#), [linux-unionfs](#), [linux-usb](#), [linux-userfs](#), [linux-watchdog](#), [linux-wireless](#), [linux-word](#), [linux-wpan](#), [linux-x11](#), [linux-x25](#), [linux-x86_64](#), [linux-xfs](#), [live-patching](#), [lm-sensors](#), [lpc-bpf](#), [lpc-netdev-bpf](#), [lpc-netdev](#), [lvs-devel](#), [mm-commits](#), [netconf2019](#), [netdev-driver-reviewers](#), [netdev](#), [netfilter-devel](#), [netfilter](#), [perfbook](#), [phone-devel](#), [platform-driver-x86](#), [rcu](#), [reiserfs-devel](#), [rust-for-linux](#), [selinux-repolicy](#), [selinux](#), [smatch](#), [sparclinux](#), [stable-commits](#), [stable-rt](#), [stable](#), [stgt](#), [target-devel](#), [testing](#), [trinity](#), [ultraio](#), [util-linux](#), [workflows](#), [xdp-newbies](#),

Listname	Number of subscribers
List: alsa-devel ; (subscribe / unsubscribe) Info: Archives: http://www.spinics.net/lists/alsa-devel/ Footer: --- To unsubscribe from this list: send the line "unsubscribe alsa-devel" in	252
List: autofs ; (subscribe / unsubscribe) Info: This is the mailing list for autofs development. Archives: Footer:	108



Patch

Contributing **units** in the kernel



Patch

Contributing units in the kernel



Developers make a **patch of development**

```
$ git format-patch -1  
--subject-prefix="PATCH, net-next"
```



```
From 0199215216978b612d4e8be11e878b87bc643033 Mon Sep 17 00:00:00 2001  
From: Juhee Kang <claudiajkang@gmail.com>  
Date: Sun, 10 Oct 2021 13:03:29 +0900  
Subject: [PATCH, net-next] mlxsw: spectrum: use netif_is_macsec() instead of  
open code  
  
Open code which is dev->priv_flags & IFF_MACSEC has already defined as  
netif_is_macsec(). So use netif_is_macsec() instead of open code.  
This patch doesn't change logic.  
  
Signed-off-by: Juhee Kang <claudiajkang@gmail.com>  
---  
include/linux/netdevice.h | 2 +-  
1 file changed, 1 insertion(+), 1 deletion(-)  
  
diff --git a/include/linux/netdevice.h b/include/linux/netdevice.h  
index 15f4a658e436..0723c1314ea2 100644  
--- a/include/linux/netdevice.h  
+++ b/include/linux/netdevice.h  
@@ -5237,7 +5237,7 @@ static inline void netif_keep_dst(struct net_device *dev)  
static inline bool netif_reduces_vlan_mtu(struct net_device *dev)  
{  
    /* TODO: reserve and use an additional IFF bit, if we get more users */  
-    return dev->priv_flags & IFF_MACSEC;  
+    return netif_is_macsec(dev);  
}  
  
extern struct pernet_operations __net_initdata loopback_net_ops;  
--  
2.34.1
```



Patch

Contributing units in the kernel



Developers make a **patch of development** and **submit it to the mailing list** by email



```
$ git format-patch -1  
--subject-prefix="PATCH, net-next"
```

```
$ git send-patch --cc "Network Maintainer  
<example@email.com>" --to "Network  
Mailing List <list@email.com>" *.patch
```



Patchwork Netdev + BPF Patches Bundles About this Login Register Mail settings

Show patches with: Submitter = Juhee Kang | 29 patches

Patch	Date	Submitter	Delegate	State
[net-next,3/3] pktgen: document the latest pktgen usage options	2021-08-25	Juhee Kang	netdev	Accepted
[net-next,2/3] samples: pktgen: add trap SIGINT for printing execution result	2021-08-25	Juhee Kang	netdev	Accepted
[net-next,1/3] samples: pktgen: fix to print when terminated normally	2021-08-25	Juhee Kang	netdev	Accepted
[net-next,2/2] samples: pktgen: add missing IPv6 option to pktgen scripts	2021-08-12	Juhee Kang	netdev	Accepted
[net-next,1/2] samples: pktgen: pass the environment variable of normal user to su	2021-08-12	Juhee Kang	netdev	Accepted
[bpf-next,v2,2/2] samples: bpf: Add the omitted xdp samples to .gitignore	2021-07-27	Juhee Kang	bpf	Accepted
[bpf-next,v2,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	2021-07-27	Juhee Kang	bpf	Accepted
[bpf-next,2/2] samples: bpf: Add the omitted xdp samples to .gitignore	2021-07-24	Juhee Kang	bpf	Superseded
[bpf-next,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	2021-07-24	Juhee Kang	bpf	Superseded



Kernel Development Process


```
$ git format-patch -1  
--subject-prefix="PATCH, net-next"
```

[PATCH net-next] mlxsw: spectrum: use netif_is_macsec() instead of open code

Open code which is dev->priv_flags & IFF_MACSEC has already defined as netif_is_macsec(). So use netif_is_macsec() instead of open code. This patch doesn't change logic.

Signed-off-by: Juhee Kang <claudiajkang@gmail.com>

Contributor create patches





Kernel Development Process


```
$ git send-patch --cc "Network Maintainer  
<example@email.com>" --to "Network Mailing List  
<list@email.com>" *.patch
```

[PATCH net-next] mlxsw: spectrum: use netif_is_macsec() instead of open code

Open code which is dev->priv_flags & IFF_MACSEC has already defined as netif_is_macsec(). So use netif_is_macsec() instead of open code. This patch doesn't change logic.

Signed-off-by: Juhee Kang <claudiajkang@gmail.com>

Contributor sends patches to the mailing list.






Kernel Development Process

**From mailing list,
review and discussion is conducted.**

```
[PATCH net-next] mlxsw: spectrum: use netif_is_macsec() instead of open code

Open code which is dev->priv_flags & IFF_MACSEC has already defined as
netif_is_macsec(). So use netif_is_macsec() instead of open code.
This patch doesn't change logic.

Signed-off-by: Juhee Kang <claudiajkang@gmail.com>
```





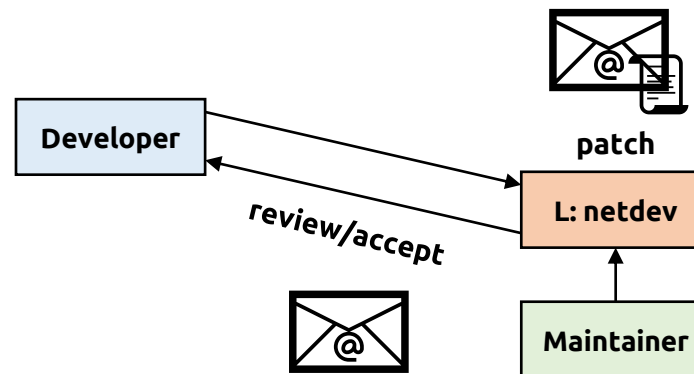
Kernel Development Process

Once the patch has been **accepted**, an email will be sent informing that it was reflected in the subsystem

```
[PATCH net-next] mlxsw: spectrum: use netif_is_macsec() instead of open code

Open code which is dev->priv_flags & IFF_MACSEC has already defined as
netif_is_macsec(). So use netif_is_macsec() instead of open code.
This patch doesn't change logic.

Signed-off-by: Juhee Kang <claudiajkang@gmail.com>
```



```
Re: [PATCH net-next] mlxsw: spectrum: use netif_is_macsec() instead of open ode


Hello:

This series was applied to netdev/net-next.git (master)
by David S. Miller <davem@davemloft.net>:
```




Kernel Development Process

Patch is reflected in net-next git tree





Kernel Development Process



index : kernel/git/torvalds/linux.git

Linux kernel source tree

about summary refs log tree **commit** diff stats

```
author Linus Torvalds <torvalds@linux-foundation.org> 2022-03-17 12:55:26 -0700
committer Linus Torvalds <torvalds@linux-foundation.org> 2022-03-17 12:55:26 -0700
commit 551acdc3c3d2b6bc97f11e31dcf960bc36343bfc (patch)
tree 1a2bf01b80da58fc613735b1c3fb03e5fb55bc2c
parent c81801eb7f2476a25d8fb27449e01b0bef46908a (diff)
parent b04683ff8f0823b869c219c78ba0d974bdddea0b5 (diff)
download linux-551acdc3c3d2b6bc97f11e31dcf960bc36343bfc.tar.gz
```


Merge tag 'net-5.17-final' of git://git.kernel.org/pub/scm/linux/kernel/git/netdev/net

Pull networking fixes from Jakub Kicinski:
"Including fixes from netfilter, ipsec, and wireless."

A few last minute revert / disable and fix patches came down from our sub-trees. We're not waiting for any fixes at this point.

Current release - regressions:

- Revert "netfilter: nat: force port remap to prevent shadowing well-known ports", restore working conntrack on asymmetric paths
- Revert "ath10k: drop beacon and probe response which leak from other channel", restore working AP and mesh mode on QCA9984
- eth: intel: fix hang during reboot/shutdown



Maintainer will Send Pull Request to Parent Tree



Kernel Development Process



index : kernel/git/torvalds/linux.git

Linux kernel source tree

about summary refs log tree **commit** diff stats

```

author Linus Torvalds <torvalds@linux-foundation.org> 2022-03-17 12:55:26 -0700
committer Linus Torvalds <torvalds@linux-foundation.org> 2022-03-17 12:55:26 -0700
commit 551acdc3c3d2b6bc97f11e31dcf960bc36343bfc (patch)
tree 1a2bf01b80da58fc613735b1c3fb03e5fb55bc2c
parent c81801eb7f2476a25d8fb27449e01b0bef46908a (diff)
parent b04683ff8f0823b869c219c78ba0d974bdddea0b5 (diff)
download linux-551acdc3c3d2b6bc97f11e31dcf960bc36343bfc.tar.gz

```

Merge tag 'net-5.17-final' of git://git.kernel.org/pub/scm/linux/kernel/git/netdev/net

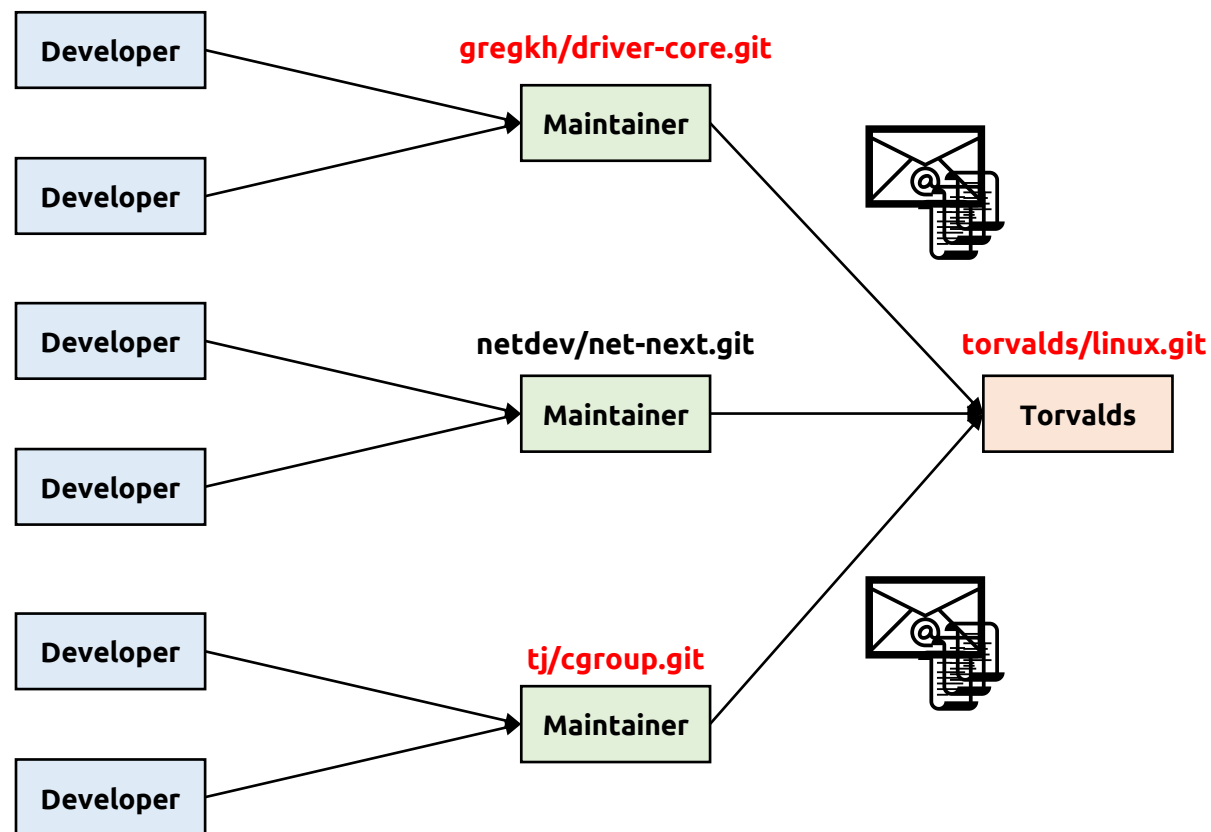
Pull networking fixes from Jakub Kicinski:
"Including fixes from netfilter, ipsec, and wireless."

A few last minute revert / disable and fix patches came down from our sub-trees. We're not waiting for any fixes at this point.

Current release - regressions:

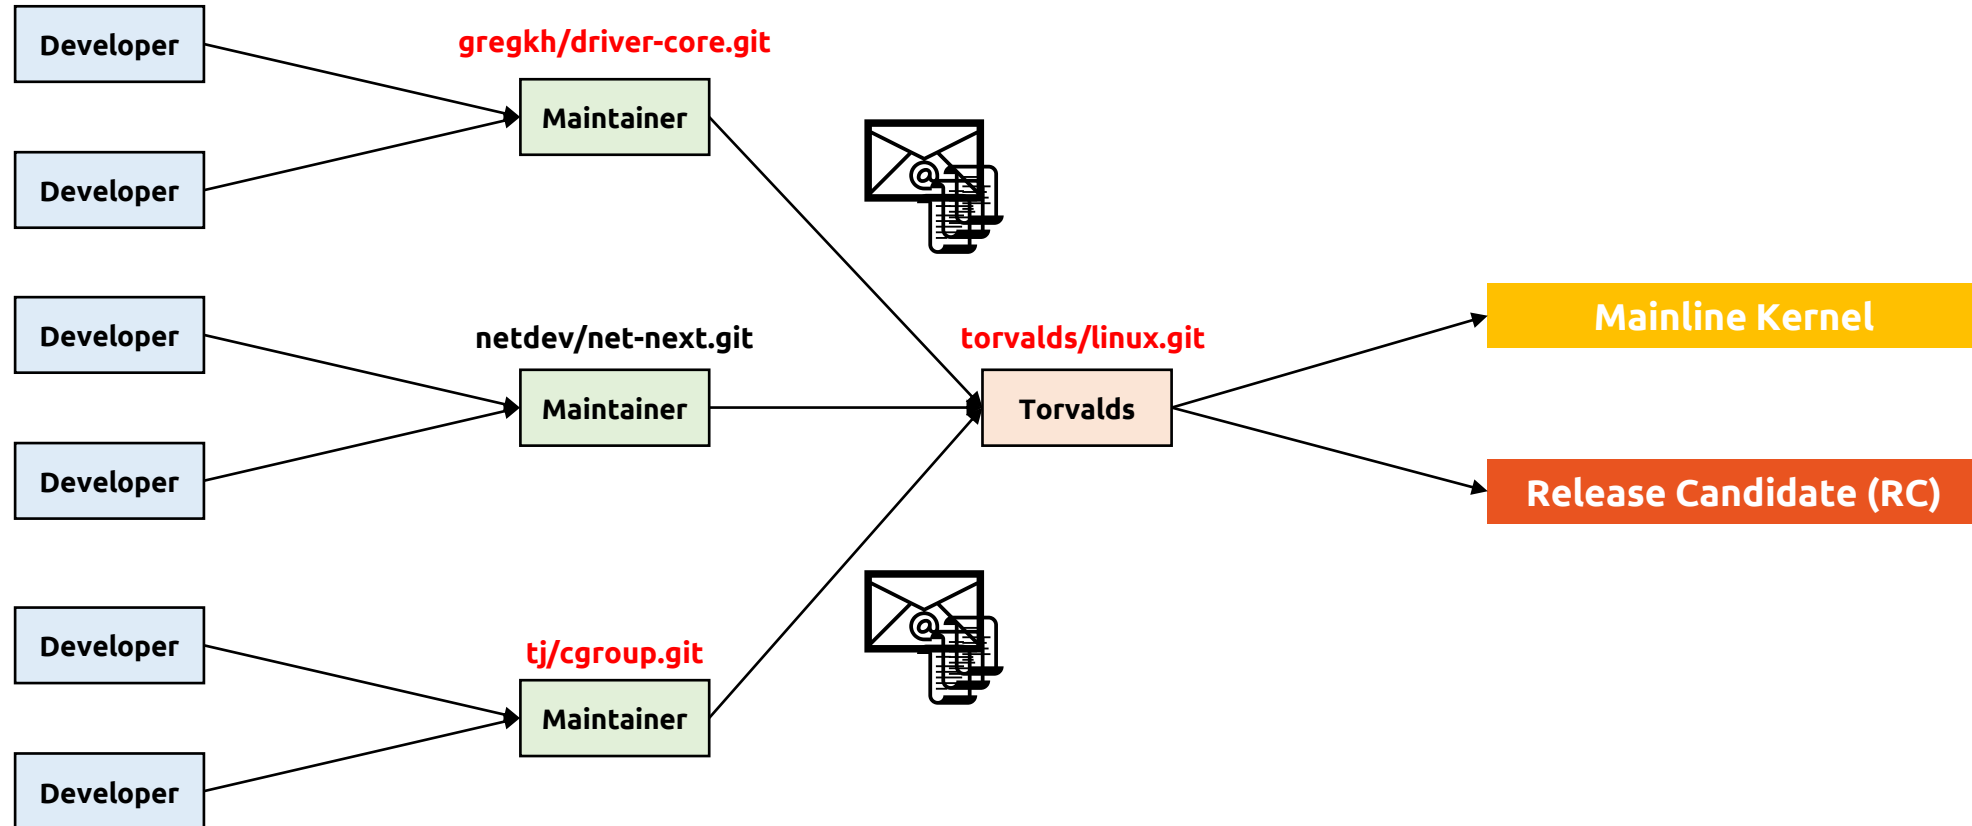
- Revert "netfilter: nat: force port remap to prevent shadowing well-known ports", restore working conntrack on asymmetric paths
- Revert "ath10k: drop beacon and probe response which leak from other channel", restore working AP and mesh mode on QCA9984
- eth: intel: fix hang during reboot/shutdown

Pull Requests are sent from various subsystems to the parent tree





Kernel Development Process



Code combined to release as new version of kernel



Linux Kernel Contribution

- ① **Documentation**
- ② **Bug-Fix**
- ③ **Feature Enhancement**
- ④ **etc..**



1. Documentation



index : kernel/git/next/linux-next.git

The linux-next integration testing tree

about summary refs log tree **commit** diff stats

```

author    Juhee Kang <claudiajkang@gmail.com> 2021-08-25 19:57:17 +0900
committer David S. Miller <davem@davemloft.net> 2021-08-25 13:44:30 +0100
commit    246b184fff1dcead37102228e3b1744ce8c9828a3 (patch)
tree      be72f30a9f1135e1c8cb152418f031bc5f9a0f0e
parent    6c882bdc4bcd63e164f05738e7677b8a62fc0e1 (diff)
download  linux-next-246b184fff1dcead37102228e3b1744ce8c9828a3.tar.gz

```

pktgen: document the latest pktgen usage options

Currently, the pktgen.rst documentation doesn't cover the latest pktgen sample usage options such as count and IPv6, and so on. Also, this documentation includes the old sample scripts which are no longer use because it was removed by the commit a4b6ade8359f ("samples/pktgen : remove remaining old pktgen sample scripts")

Thus, this commit documents pktgen sample usage using the latest options and removes old sample scripts, and fixes a minor typo.

Signed-off-by: Juhee Kang <claudiajkang@gmail.com>
Signed-off-by: David S. Miller <davem@davemloft.net>

Diffstat

```
-rw-r--r-- Documentation/networking/pktgen.rst 18
```

1 files changed, 8 insertions, 10 deletions

```

diff --git a/Documentation/networking/pktgen.rst b/Documentation/networking/pktgen.rst
index 7afa1c9f11838..1225f0f63ff07 100644
--- a/Documentation/networking/pktgen.rst
+++ b/Documentation/networking/pktgen.rst
@@ -248,26 +248,24 @@ Usage:::

```

```

-i : ($DEV)      output interface/device (required)
-s : ($PKT_SIZE) packet size
- -d : ($DEST_IP) destination IP
+ -d : ($DEST_IP) destination IP. CIDR (e.g. 198.18.0.0/15) is also allowed
-m : ($DST_MAC) destination MAC-addr
+ -p : ($DST_PORT) destination PORT range (e.g. 433-444) is also allowed
-t : ($THREADS) threads to start
+ -f : ($F_THREAD) index of first thread (zero indexed CPU number)
- -c : ($SKB_CLONE) SKB clones send before alloc new SKB
+ -n : ($COUNT)  num messages to send per thread, 0 means indefinitely
- -b : ($BURST)   HW level bursting of SKBs
- -v : ($VERBOSE) verbose
- -x : ($DEBUG)   debug
+ -6 : ($IPv6)    IPv6
+ -w : ($DELAY)   Tx Delay value (ns)
+ -a : ($APPEND)  Script will not reset generator's state, but will append its config

```

Improve the documentation

- add undocumented or elaborate details



2. Bug Fix



index : kernel/git/next/linux-next.git

The linux-next integration testing tree

[about](#) [summary](#) [refs](#) [log](#) [tree](#) **[commit](#)** [diff](#) [stats](#)

```

author    Juhee Kang <claudiajkang@gmail.com> 2021-10-04 21:14:38 +0900
committer Pablo Neira Ayuso <pablo@netfilter.org> 2021-10-07 19:35:57 +0200
commit    902c0b1887522a099aa4e1e6b4b476c2fe5dd13e (patch)
tree      37a3ee3d04c0a7d451b3d6c30f1794bb6d266276
parent    8d6c414cd2fb74aa6812e9bfec6178f8246c4f3a (diff)
download  linux-next-902c0b1887522a099aa4e1e6b4b476c2fe5dd13e.tar.gz

```

netfilter: xt_IDLETIMER: fix panic that occurs when timer_type has garbage value

Currently, when the rule related to IDLETIMER is added, idletimer_tg timer structure is initialized by kmalloc on executing idletimer_tg_create function. However, in this process timer->timer_type is not defined to a specific value. Thus, timer->timer_type has garbage value and it occurs kernel panic. So, this commit fixes the panic by initializing timer->timer_type using kzalloc instead of kmalloc.

```

Test commands:
# iptables -A OUTPUT -j IDLETIMER --timeout 1 --label test
$ cat /sys/class/xt_idletimer/timers/test
Killed

```

```

Splat looks like:
BUG: KASAN: user-memory-access in alarm_expires_remaining+0x49/0x70
Read of size 8 at addr 0000002e8c7bc4c8 by task cat/917
CPU: 12 PID: 917 Comm: cat Not tainted 5.14.0+ #3 79940a339f71eb14fc81aee1757a20d5bf13eb0e
Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS 1.13.0-lubuntu1.1 04/01/2014
Call Trace:
 dump_stack_lvl+0x6e/0x9c
 kasan_report.cold+0x112/0x117
 ? alarm_expires_remaining+0x49/0x70
 __asan_load8+0x86/0xb0
 alarm_expires_remaining+0x49/0x70
 idletimer_tg_show+0xe5/0x19b [xt_IDLETIMER 11219304af9316a21bee5ba9d58f76a6b9bcc6d]
 dev_attr_show+0x3c/0x60
 sysfs_kf_seq_show+0x11d/0x1f0
 ? device_remove_bin_file+0x20/0x20

```

Fixing Bugs

- fix code to work as expected



3. Feature Enhancement



index : kernel/git/next/linux-next.git

The linux-next integration testing tree

about summary refs log tree **commit** diff stats

author Juhee Kang <claudiajkang@gmail.com> 2021-08-13 00:08:13 +0900
 committer David S. Miller <davem@davemloft.net> 2021-08-16 11:02:09 +0100
 commit 0f0c4f1b72e090b23131700bb155944cc28b2a7b (patch)
 tree 1dc881a9d04a4256e1f91415a6e7f9ec0796b729
 parent 7caeabd726f22e6a6c44c434574fb489986e5baa (diff)
 download linux-next-0f0c4f1b72e090b23131700bb155944cc28b2a7b.tar.gz

samples: pktgen: add missing IPv6 option to pktgen scripts

Currently, "sample04" and "sample05" are not working properly when running with an IPv6 option ("-6"). The commit 0f06a6787e05 ("samples: Add an IPv6 "-6" option to the pktgen scripts") has omitted the addition of this option at "sample04" and "sample05".

In order to support IPv6 option, this commit adds logic related to IPv6 option.

Fixes: 0f06a6787e05 ("samples: Add an IPv6 "-6" option to the pktgen scripts")

Signed-off-by: Juhee Kang <claudiajkang@gmail.com>
Signed-off-by: David S. Miller <davem@davemloft.net>

Diffstat

```
-rwxr-xr-x samples/pktgen/pktgen_sample04_many_flows.sh 12
-rwxr-xr-x samples/pktgen/pktgen_sample05_flow_per_thread.sh 12
```

2 files changed, 14 insertions, 10 deletions

```
diff --git a/samples/pktgen/pktgen_sample04_many_flows.sh b/samples/pktgen/pktgen_sample04_many_flows.sh
index 56c5f5af350f6..cf151f861506d 100755
--- a/samples/pktgen/pktgen_sample04_many_flows.sh
+++ b/samples/pktgen/pktgen_sample04_many_flows.sh
@@ -13,13 +13,15 @@ root_check_run_with_sudo "$@"
 # Parameter parsing via include
 source ${basedir}/parameters.sh
 # Set some default params, if they didn't get set
-[ -z "$DEST_IP" ]    && DEST_IP="198.18.0.42"
+if [ -z "$DEST_IP" ]; then
+  [ -z "$IP6" ] && DEST_IP="198.18.0.42" || DEST_IP="f000::1"
+fi
[ -z "$DST_MAC" ]    && DST_MAC="90:e2:ba:ff:ff:ff"
[ -z "$CLONE_SKB" ]  && CLONE_SKB="0"
[ -z "$COUNT" ]    && COUNT="0" # Zero means indefinitely
if [ -n "$DEST_IP" ]; then
```

Adding a new feature

- for improved usability



3. Feature Enhancement



index : kernel/git/next/linux-next.git

The linux-next integration testing tree

about summary refs log tree **commit** diff stats

```

author    Juhee Kang <claudiajkang@gmail.com> 2022-02-05 15:40:38 +0000
committer David S. Miller <davem@davemloft.net> 2022-02-06 10:55:52 +0000
commit    4acc45db711586dadb98e681ccb93f29372907a2 (patch)
tree      f69f105c0d7561d7caacec85ef42c2d9463843bf
parent    5a8fb33e530512ee67a11b90f3451a4f030f4b01 (diff)
download  linux-next-4acc45db711586dadb98e681ccb93f29372907a2.tar.gz

```

net: hsr: use hlist_head instead of list_head for mac addresses

Currently, HSR manages mac addresses of known HSR nodes by using list_head. It takes a lot of time when there are a lot of registered nodes due to finding specific mac address nodes by using linear search. We can be reducing the time by using hlist. Thus, this patch moves list_head to hlist_head for mac addresses and this allows for further improvement of network performance.

```

Condition: registered 10,000 known HSR nodes
Before:
# iperf3 -c 192.168.10.1 -i 1 -t 10
Connecting to host 192.168.10.1, port 5201
[ 5] local 192.168.10.2 port 59442 connected to 192.168.10.1 port 5201
[ ID] Interval      Transfer    Bitrate    Retr  Cwnd
[ 5] 0.00-1.49 sec  3.75 MBytes 21.1 Mbits/sec  0  158 KBytes
[ 5] 1.49-2.05 sec  1.25 MBytes 18.7 Mbits/sec  0  166 KBytes
[ 5] 2.05-3.06 sec  2.44 MBytes 20.3 Mbits/sec  56  16.9 KBytes
[ 5] 3.06-4.08 sec  1.43 MBytes 11.7 Mbits/sec  11  38.0 KBytes
[ 5] 4.08-5.00 sec   951 KBytes  8.49 Mbits/sec  0  56.3 KBytes

After:
# iperf3 -c 192.168.10.1 -i 1 -t 10
Connecting to host 192.168.10.1, port 5201
[ 5] local 192.168.10.2 port 36460 connected to 192.168.10.1 port 5201
[ ID] Interval      Transfer    Bitrate    Retr  Cwnd
[ 5] 0.00-1.00 sec  7.39 MBytes 62.0 Mbits/sec  3  130 KBytes
[ 5] 1.00-2.00 sec  5.06 MBytes 42.4 Mbits/sec  16  113 KBytes
[ 5] 2.00-3.00 sec  8.58 MBytes 72.0 Mbits/sec  42  94.3 KBytes
[ 5] 3.00-4.00 sec  7.44 MBytes 62.4 Mbits/sec  2  131 KBytes
[ 5] 4.00-5.07 sec  8.13 MBytes 63.5 Mbits/sec  38  92.9 KBytes

```

Signed-off-by: Juhee Kang <claudiajkang@gmail.com>
Signed-off-by: David S. Miller <davem@davemloft.net>

Diffstat

```

-rw-r--r-- net/hsr/hsr_debugfs.c 40

```

Enhance feature

- performance improvement



Now let's do Contributions to Linux Kernel



3

Step by Step Kernel Contribution



Kernel Patch in Action

In this scenario, suppose we fix a bug.



Discover a Bug

```
$ cd samples/bpf
$ make -j
$ ./tracex7
sh: 1: Syntax error: word unexpected (expecting ")")
```

An **error** occurred
while running the kernel test script.



Configure Environment: kernel development



Fedora
Debian
OpenSUSE
Ubuntu

Prepare Linux in VM



Configure Environment: kernel development

```
# apt-get install build-essential vim git cscope  
libncurses-dev libssl-dev bison flex bison
```

Install development packages on Linux VM



Configure Environment: kernel development

```
# apt-get install build-essential vim git cscope
libncurses-dev libssl-dev bison flex bison
# systemctl enable serial-getty@ttyS0.service
# systemctl start serial-getty@ttyS0.service
# systemctl daemon-reload
# vim /etc/default/grub          GRUB_DEFAULT=3
                                GRUB_TIMEOUT_STYLE=menu
                                GRUB_TIMEOUT=3
# update-grub2
```

Configure Linux for a debugging



Configure Environment: get kernel source

```
$ git clone https://git.kernel.org/pub/scm/linux/kernel/git/netdev/net-next.git
$ cd net-next
$ git checkout -f main
```



index : kernel/git/netdev/net-next.git

Netdev Group's -next networking tree

master

Netdev Group

[about](#) [summary](#) [refs](#) [log](#) [tree](#) [commit](#) [diff](#) [stats](#)

Branch	Commit message	Author	Age
main	Merge tag 'net-next-6.1' of git://git.kernel.org/pub/scm/linux/kernel/git/net...	Linus Torvalds	3 days
master	Merge tag 'net-next-6.1' of git://git.kernel.org/pub/scm/linux/kernel/git/net...	Linus Torvalds	3 days

Tag	Download	Author	Age
v6.0	net-next-6.0.tar.gz	Linus Torvalds	5 days
v6.0-rc7	net-next-6.0-rc7.tar.gz	Linus Torvalds	12 days
v6.0-rc6	net-next-6.0-rc6.tar.gz	Linus Torvalds	3 weeks
v6.0-rc5	net-next-6.0-rc5.tar.gz	Linus Torvalds	4 weeks
v6.0-rc4	net-next-6.0-rc4.tar.gz	Linus Torvalds	5 weeks
v6.0-rc3	net-next-6.0-rc3.tar.gz	Linus Torvalds	6 weeks
v6.0-rc2	net-next-6.0-rc2.tar.gz	Linus Torvalds	7 weeks
v6.0-rc1	net-next-6.0-rc1.tar.gz	Linus Torvalds	8 weeks
v5.19	net-next-5.19.tar.gz	Linus Torvalds	2 months
v5.19-rc8	net-next-5.19-rc8.tar.gz	Linus Torvalds	2 months
[...]			



Reproduce Bug: find the cause

```
$ cd samples/bpf
$ make -j
$ ./tracex7
sh: 1: Syntax error: word unexpected (expecting ")")
```

An **error** occurred
while running the kernel test script.

Maybe Kernel test script has a Bug.



Reproduce Bug: find the cause

```
$ cd samples/bpf
$ make -j
$ ./tracex7
sh: 1: Syntax error: word unexpected (expecting ")")
# vim tracex7_user.c
```

An **error** occurred while running the kernel test script.

```
1 tracex7_user.c +
int main(int argc, char **argv)
{
    // ... SNIP ...
    char filename[256];
    char command[256];
    int ret = 0;
    FILE *f;

    snprintf(filename, sizeof(filename), "%s_kern.o", argv[0]);
    ++-- 24 lines: obj = bpf_object__open_file(filename, NULL);-----
    snprintf(command, 256, "mount %s tmpmnt/", argv[1]);
    f = popen(command, "r");
    ret = pclose(f);

cleanup:
    bpf_link__destroy(link);
    bpf_object__close(obj);
    return ret ? 0 : 1;
}
NORMAL tracex7_user.c + 18% 8:0
```

Problem caused by **failure to check argument at runtime**



Fix source code: resolve problem

```
$ git diff
```

```
diff --git a/samples/bpf/test_override_return.sh b/samples/bpf/test_override_return.sh
index e68b9ee6814b8b..35db26f736b9d9 100755
--- a samples/bpf/test\_override\_return.sh
+++ b samples/bpf/test\_override\_return.sh
@@ -1,5 +1,6 @@
#!/bin/bash
+rm -r tmpmnt
rm -f testfile.img
dd if=/dev/zero of=testfile.img bs=1M seek=1000 count=1
DEVICE=$(losetup --show -f testfile.img)
diff --git a/samples/bpf/tracex7_user.c b/samples/bpf/tracex7_user.c
index fdcd6580dd736a..8be7ce18d3ba05 100644
    samples/bpf/tracex7\_user.c
    samples/bpf/tracex7\_user.c
@@ -14,6 +14,11 @@ int main(int argc, char **argv)
int ret = 0;
FILE *f;
+ if (!argv[1]) {
+ fprintf(stderr, "ERROR: Run with the btrfs device argument!\n");
+ return 0;
+ }
+
snprintf(filename, sizeof(filename), "%s_kern.o", argv[0]);
```



Write a commit message

After modifying the source code, **create a commit** with detailed message

```
$ git commit -s
```



Write a commit message

Having a trouble with writing commit message? Check previous commit log.

< change file name >

```
$ git log -- samples/bpf/tracex7_user.c
```

```
commit 2bf3e2ef425bc2a164f10b554b7db6a8b4090ef4
```

```
samples: bpf: include bpf/bpf.h instead of local libbpf.h
```

There are two files in the tree called libbpf.h which is becoming problematic. Most samples don't actually need the local libbpf.h they simply include it to get to bpf/bpf.h. Include bpf/bpf.h directly instead.

...

```
commit 965de87e54b803223bff703ea6b2a76c056695ae
```

```
Author: Josef Bacik <jbacik@fb.com>
```

```
Date: Mon Dec 11 11:36:49 2017 -0500
```

```
samples/bpf: add a test for bpf_override_return
```



Write a commit message

```
samples: bpf: Fix tracex7 error raised on the missing argument
```

The current behavior of 'tracex7' doesn't consist with other bpf samples tracex{1..6}. Other samples do not require any argument to run with, but tracex7 should be run with btrfs device argument. (it should be executed with test_override_return.sh)

Currently, tracex7 doesn't have any description about how to run this program and raises an unexpected error. And this result might be confusing since users might not have a hunch about how to run this program.

. . .

In order to fix this error, this commit adds logic to report a message and exit when running this program with a missing argument.

Additionally in test_override_return.sh, there is a problem with multiple directory(tmpmnt) creation. So in this commit adds a line with removing the directory with every execution.

```
Signed-off-by: Juhee Kang <claudiajkang@gmail.com>
```

Write commit message in detail.

```
samples: bpf: Fix tracex7 error raised on the missing argument
```

The current behavior of 'tracex7' doesn't consist with other bpf samples tracex{1..6}. Other samples do not require any argument to run with, but tracex7 should be run with btrfs device argument. (it should be executed with test_override_return.sh)

Currently, tracex7 doesn't have any description about how to run this program and raises an unexpected error. And this result might be confusing since users might not have a hunch about how to run this program.

```
// Current behavior
# ./tracex7
sh: 1: Syntax error: word unexpected (expecting ")")
// Fixed behavior
# ./tracex7 ERROR: Run with the btrfs device argument!
```

In order to fix this error, this commit adds logic to report a message and exit when running this program with a missing argument.

Additionally in test_override_return.sh, there is a problem with multiple directory(tmpmnt) creation. So in this commit adds a line with removing the directory with every execution.

```
Signed-off-by: Juhee Kang <claudiajkang@gmail.com>
```

Commit Header
- less than 50 characters

Commit description

- elaborate detail
- 72 characters per line

**If necessary,
describe how to reproduce
(for detailed explanation)**

Commit Tag

- Add Signed-off-by Tag



Other Commit Tags

```
mlxsw: spectrum: use netif_is_macsec() instead of open code
=== Commit Body ===
Signed-off-by: Juhee Kang <claudiajkang@gmail.com>
```

```
samples: bpf: Fix trace7 error raised on the missing argument
=== Commit Body ===
Acked-by: Yonghong Song <yhs@fb.com>
```

```
amt: Use BIT macros instead of open codes
=== Commit Body ===
Reviewed-by: Taehee Yoo <ap420073@gmail.com>
Tested-by: Taehee Yoo <ap420073@gmail.com>
```

```
net: hsr: fix hsr build error when lockdep is not enabled
=== Commit Body ===
Reported-by: Eric Dumazet <eric.dumazet@gmail.com>
```

```
bpf_glue: include errno.h
=== Commit Body ===
Fixes: ac4e0913beb1 ("bpf: Export bpf syscall wrapper")
```

Commit Tag Detail

1) Signed-off-by

understood the licensing policy, agree to submit commit

2) Acked-by

maintainer reviewed the patch and decided to approve the patch

3) Reviewed-by

Reviewer or another developer reviewed the code for the patch

4) Tested-by

Reviewer or other developer tested for changes to the patch

5) Reported-by / Suggested-by

Tags left to indicate who found the bug or suggested the feature

6) Fixes

Tags that track bugs(commit) that this patch fixes



Configure Environment: prepare to send patch

```
# apt-get install git-email  
  
$ git config --global sendemail.smtpencryption tls  
$ git config --global sendemail.smtpserver <mailserver>  
$ git config --global sendemail.smtpuser <email>  
$ git config --global sendemail.smtpserverport 587  
$ git config --global sendemail.smtppass <password>
```

send patches with **git-email**

Configure **smtp account** with git config
ex) gmail smtp



Create patch and check format

```
$ git format-patch -1 --subject-prefix="PATCH, bpf-next"
```



Create patch and check format

<commit count> **<Add kernel tree to mail prefix>**

```
$ git format-patch -1 --subject-prefix="PATCH, bpf-next"
```



Create patch and check format

<commit count> **<Add kernel tree to mail prefix>**

```
$ git format-patch -1 --subject-prefix="PATCH, bpf-next"  
001-samples-bpf-Fix-tracex7-error-raised-on-the-missing-.patch
```



Create patch and check format

```
$ git format-patch -1 --subject-prefix="PATCH, bpf-next"  
0001-samples-bpf-Fix-tracex7-error-raised-on-the-missing-.patch  
  
$ apt-get install codespell python-ply python-git  
$ ./scripts/checkpatch.pl --strict --codespell 0001-*.patch  
                <created patch file here>
```



Create patch and check format

```
$ git format-patch -1 --subject-prefix="PATCH, bpf-next"  
0001-samples-bpf-Fix-tracex7-error-raised-on-the-missing-.patch  
  
$ apt-get install codespell python-ply python-git  
$ ./scripts/checkpatch.pl --strict --codespell 0001-*.patch  
total: 0 errors, 0 warnings, 0 checks, 17 lines checked  
  
0001-*.patch has no obvious style problems and is ready for submission.  
    < Good to go! >
```



Find patch mail recipients and send

<created patch file here>

```
$ ./scripts/get_maintainer.pl --no-multiline --no-rolestats 0001-*.patch
```

**Find maintainers
& mailing lists**



Find patch mail recipients and send

```
$ ./scripts/get_maintainer.pl --no-multiline --no-rolestats 0001-*.patch  
Alexei Starovoitov <ast@kernel.org>, Daniel Borkmann <daniel@iogearbox.net>,  
Andrii Nakryiko <andrii@kernel.org>, netdev@vger.kernel.org, bpf@vger.kernel.org
```

**Find maintainers
& mailing lists**



Find patch mail recipients and send

```
$ ./scripts/get_maintainer.pl --no-multiline --no-rolestats 0001-*.patch  
Alexei Starovoitov <ast@kernel.org>, Daniel Borkmann <daniel@iogearbox.net>,  
Andrii Nakryiko <andrii@kernel.org>, netdev@vger.kernel.org, bpf@vger.kernel.org
```

```
$ git send-email --to="maintainer mail" --to="mailing list mail" 0001-*.patch  
Send this email? ([y]es|[n]o|[q]uit|[a]ll): Y
```

```
[...]
```

```
Result: OK
```

Send mail



Make sure mail has been sent

<input type="checkbox"/>	★	▶	me	Inbox	[PATCH net-next 2/2] samples: pktgen: add missing IPv6 option to pktgen scripts - Currently, "sample04" and "sample05" are not working properly when running with an IPv...	8/13/21
<input type="checkbox"/>	★	▶	me	Inbox	[PATCH net-next 1/2] samples: pktgen: pass the environment variable of normal user to sudo - All pktgen samples can use the environment variable instead of option param...	8/13/21
<input type="checkbox"/>	★	▶	me	Inbox	[PATCH net-next 0/2] samples: pktgen: enhance the usability of pktgen samples - This patchset improves the usability of pktgen samples by adding an option for propagati...	8/13/21
<input type="checkbox"/>	☆	▶	Majordomo 2	Inbox	Confirmation for subscribe netdev - -- Someone (possibly you) has requested that your email address be added to or deleted from the mailing list "netdev@vger.kernel.org"....	8/5/21
<input type="checkbox"/>	★	▶	me .. patchwork-b. 6	Inbox	[bpf-next 1/2] samples: bpf: Fix tracex7 error raised on the missing argument - Hello: This series was applied to bpf/bpf-next.git (refs/heads/master): On Tue, 27 Jul 2021 04:1...	7/28/21
<input type="checkbox"/>	☆	▶	me	Inbox	[bpf-next v2 2/2] samples: bpf: Add the omitted xdp samples to .gitignore - There are recently added xdp samples (xdp_redirect_map_multi and xdpsock_ctrl_proc) which ar...	7/27/21
<input type="checkbox"/>	★	▶	me, Yonghong 2	Inbox	[bpf-next 2/2] samples: bpf: Add the omitted xdp samples to .gitignore - On 7/24/21 8:21 AM, Juhee Kang wrote: > There are recently added xdp samples (xdp_redirect_map_...	7/26/21

You can check the patch through **Gmail**



Make sure mail has been sent

Show patches with: Submitter = Juhee Kang | 29 patches

Patch	Series	A/R/T	S/W/F	▲ Date	Submitter	Delegate	State
[net-next,3/3] pktgen: document the latest pktgen usage options	samples: pktgen: enhance the ability to print the execution results of samples	---	14 1 -	2021-08-25	Juhee Kang	netdev	Accepted
[net-next,2/3] samples: pktgen: add trap SIGINT for printing execution result	samples: pktgen: enhance the ability to print the execution results of samples	---	14 1 -	2021-08-25	Juhee Kang	netdev	Accepted
[net-next,1/3] samples: pktgen: fix to print when terminated normally	samples: pktgen: enhance the ability to print the execution results of samples	---	14 1 -	2021-08-25	Juhee Kang	netdev	Accepted
[net-next,2/2] samples: pktgen: add missing IPv6 option to pktgen scripts	samples: pktgen: enhance the usability of pktgen samples	---	13 - 2	2021-08-12	Juhee Kang	netdev	Accepted
[net-next,1/2] samples: pktgen: pass the environment variable of normal user to sudo	samples: pktgen: enhance the usability of pktgen samples	---	14 1 -	2021-08-12	Juhee Kang	netdev	Accepted
[bpf-next,v2,2/2] samples: bpf: Add the omitted xdp samples to .gitignore	[bpf-next,v2,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	---	14 1 -	2021-07-27	Juhee Kang	bpf	Accepted
[bpf-next,v2,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	[bpf-next,v2,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	1 --	14 1 -	2021-07-27	Juhee Kang	bpf	Accepted
[bpf-next,2/2] samples: bpf: Add the omitted xdp samples to .gitignore	[bpf-next,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	1 --	14 1 -	2021-07-24	Juhee Kang	bpf	Superseded
[bpf-next,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	[bpf-next,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	1 --	14 1 -	2021-07-24	Juhee Kang	bpf	Superseded

<https://patchwork.kernel.org/project/netdevbpf/list>

You can check the patch through **Patchwork** as well
(Patchwork : Discussions and development of subsystem View patches.)



Check for review

```
From: Andrii Nakryiko <andrii.nakryiko@gmail.com>  
To: Juhee Kang <claudiajkang@gmail.com>  
Cc: Daniel Borkmann <daniel@iogearbox.net>, Alexei Starovoitov <ast@kernel.org>,  
    Andrii Nakryiko <andrii@kernel.org>, Networking <netdev@vger.kernel.org>, bpf <bpf@vger.kernel.org>  
Subject: Re: [bpf-next 1/2] samples: bpf: Fix tracex7 error raised on the missing argument
```

```
On Sat, Jul 24, 2021 at 8:21 AM Juhee Kang <claudiajkang@gmail.com> wrote:
```

```
[...]  
> diff --git a/samples/bpf/test_override_return.sh b/samples/bpf/test_override_return.sh  
> index e68b9ee6814b..6480b55502c7 100755  
> --- a/samples/bpf/test_override_return.sh  
> +++ b/samples/bpf/test_override_return.sh  
> @@ -1,5 +1,6 @@  
> #!/bin/bash  
>  
> +rm -rf tmpmnt
```

```
Do we need -rf or -r would do?
```

```
[...]  
> --  
> 2.27.0
```



Reply to review

```
From: Juhee Kang <claudiajkang@gmail.com>
To: Andrii Nakryiko <andrii.nakryiko@gmail.com>
Cc: Daniel Borkmann <daniel@iogearbox.net>, Alexei Starovoitov <ast@kernel.org>,
    Andrii Nakryiko <andrii@kernel.org>, Networking <netdev@vger.kernel.org>, bpf <bpf@vger.kernel.org>
Subject: Re: [bpf-next 1/2] samples: bpf: Fix tracex7 error raised on the missing argument
```

```
On Tue, Jul 27, 2021 at 5:08 AM Andrii Nakryiko <andrii.nakryiko@gmail.com> wrote:
```

```
> On Sat, Jul 24, 2021 at 8:21 AM Juhee Kang <claudiajkang@gmail.com> wrote:
```

```
[...]
```

```
>> +rm -rf tmpmnt
```

```
> Do we need -rf or -r would do?
```

```
It works properly using only -r.
```

```
Thanks for pointing it out! I will stick to this method.
```

```
I will send the next version as soon as possible.
```

```
[...]
```

```
>> 2.27.0
```



Modify Code and create second Patch

< Specify version of the patch >

```
$ git format-patch -1 -v2 --subject-prefix="PATCH, bpf-next"  
v2-0001-samples-bpf-Fix-tracex7-error-raised-on-the-missing-.patch
```

```
$ ./scripts/checkpatch.pl --strict --codespell v2-0001-*.patch  
total: 0 errors, 0 warnings, 0 checks, 17 lines checked
```

v2-0001-*.patch has **no obvious style problems** and is ready for submission.



Find patch mail recipients and send

```
$ ./scripts/get_maintainer.pl --no-multiline --no-rolestats v2-0001-*.patch  
Alexei Starovoitov <ast@kernel.org>, Daniel Borkmann <daniel@iogearbox.net>,  
Andrii Nakryiko <andrii@kernel.org>, netdev@vger.kernel.org, bpf@vger.kernel.org
```

```
$ git send-email --to="maintainer mail" --to="mailing list mail" v2-0001-*.patch  
Send this email? ([y]es|[n]o|[q]uit|[a]ll): Y
```

[...]

Result: OK

Send mail



Patch accept

```
From: patchwork-bot+netdevbpf <patchwork-bot+netdevbpf@kernel.org>  
To: Juhee Kang <claudiajkang@gmail.com>  
Subject: Re: [bpf-next v2 1/2] samples: bpf: Fix tracex7 error raised on the missing argument
```

Hello:

This series was applied to [bpf/bpf-next.git](https://git.kernel.org/bpf/bpf-next.git) (refs/heads/master):

On Tue, 27 Jul 2021 04:10:55 +0000 you wrote:

```
> The current behavior of 'tracex7' doesn't consist with other bpf samples  
> tracex{1..6}. Other samples do not require any argument to run with, but  
> tracex7 should be run with btrfs device argument. (it should be executed  
> with test_override_return.sh)  
> [...]
```

Here is the summary with links:

- [bpf-next,v2,1/2] samples: bpf: Fix tracex7 error raised on the missing argument
<https://git.kernel.org/bpf/bpf-next/c/7d07006f0592>
- [bpf-next,v2,2/2] samples: bpf: Add the omitted xdp samples to .gitignore
<https://git.kernel.org/bpf/bpf-next/c/05e9b4f60d31>

You are awesome, thank you!

You can also check it through [Gmail](#)



Patch accept

Patchwork Netdev + BPF Patches Bundles About this Login Register Mail settings

Show patches with: Submitter = Juhee Kang | 29 patches

Patch	Date	Submitter	Delegate	State
[net-next,3/3] pktgen: document the latest pktgen usage options	2021-08-25	Juhee Kang	netdev	Accepted
[net-next,2/3] samples: pktgen: add trap SIGINT for printing execution result	2021-08-25	Juhee Kang	netdev	Accepted
[net-next,1/3] samples: pktgen: fix to print when terminated normally	2021-08-25	Juhee Kang	netdev	Accepted
[net-next,2/2] samples: pktgen: add missing IPv6 option to pktgen scripts	2021-08-12	Juhee Kang	netdev	Accepted
[net-next,1/2] samples: pktgen: pass the environment variable of normal user to su	2021-08-12	Juhee Kang	netdev	Accepted
[bpf-next,v2,2/2] samples: bpf: Add the omitted xdp samples to .gitignore	2021-07-27	Juhee Kang	bpf	Accepted
[bpf-next,v2,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	2021-07-27	Juhee Kang	bpf	Accepted
[bpf-next,2/2] samples: bpf: Add the omitted xdp samples to .gitignore	2021-07-24	Juhee Kang	bpf	Superseded
[bpf-next,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	2021-07-24	Juhee Kang	bpf	Superseded

<https://patchwork.kernel.org/project/netdevbpf/list>

If the patch is reflected, **the state** about the patch in **patchwork** changes to **accept**.



This was a quick guide to Linux kernel contribution



Need help with kernel patching?



Google Summer of Code

Participate in the kernel-related Contribution program



Need help with kernel patching in Korea?


**2022 오픈소스
컨트리뷰션 아카데미** Open Source
Contribution
Academy

기여(Contribution) + 마라톤(Marathon) = 컨트리뷰톤!

contribute

주최 과학기술정보통신부 주관 nipa 정보통신산업진흥원 OPENUP

Participate in the kernel-related Contribution program



Hope you've overcome the barriers to kernel contribution!



Q&A

